

Sun Netra T5440 Server

Administration Guide



Part No. 820-4443-11
April 2010, Revision A

Copyright © 2008, 2010 Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2008, 2010, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.



Adobe PostScript

Contents

Preface ix

1. Configuring the System Console 1

Communicating With the System 1

What the System Console Does 3

Using the System Console 3

Default System Console Connection Through the Serial Management and Network Management Ports 4

Alternative System Console Configuration 6

Accessing the System Console Through a Graphics Monitor 7

Accessing the Service Processor 7

Using the Serial Management Port 7

▼ To Use the Serial Management Port 7

Activating the Network Management Port 8

▼ To Activate the Network Management Port 9

Accessing the System Console Through a Terminal Server 10

▼ To Access the System Console Through a Terminal Server 10

Accessing the System Console Through a Tip Connection 12

▼ To Access the System Console Through the Tip Connection 13

Modifying the `/etc/remote` File 14

▼ To Modify the <code>/etc/remote</code> File	14
Accessing the System Console Through an Alphanumeric Terminal	15
▼ To Access the System Console Through an Alphanumeric Terminal	16
Accessing the System Console Through a Local Graphics Monitor	16
▼ To Access the System Console Through a Local Graphics Monitor	17
Switching Between the Service Processor and the System Console	18
▼ To Access the System Console Through a Local Graphics Monitor	20
ILOM -> Prompt	20
Access Through Multiple Controller Sessions	21
Reaching the -> Prompt	21
OpenBoot ok Prompt	22
OpenBoot™ ok Prompt Not Available After the Solaris OS Starts	23
Reaching the ok Prompt	23
Graceful Shutdown	23
▼ To Use ILOM set <code>/HOST send_break_action=break, start /SP/console</code> Commands, or Break Key	24
Manual System Reset	25
Graceful Reset of the Control Domain With ILOM reset Command	25
▼ To Obtain the ok Prompt	25
For More Information	26
System Console OpenBoot Configuration Variable Settings	26
2. Managing RAS Features and System Firmware	29
ILOM and the Service Processor	30
Logging In To ILOM	30
▼ To Log In To ILOM	30
▼ To View System Fault Information	31
Status Indicators	31
Interpreting System LEDs	32

Bezel Server Status Indicators	33
Alarm Status Indicators	35
Controlling the Locator LED	37
▼ To Control the Locator LED	37
OpenBoot Emergency Procedures	38
OpenBoot Emergency Procedures for the Sun Netra T5440 System	38
Stop-N Functionality	38
▼ To Restore OpenBoot Configuration Defaults	39
Stop-F Functionality	40
Stop-D Functionality	40
Automatic System Recovery	40
Auto-Boot Options	41
Error Handling Summary	41
Reset Scenarios	42
Automatic System Recovery User Commands	43
Enabling and Disabling Automatic System Recovery	43
▼ To Enable Automatic System Recovery	44
▼ To Disable Automatic System Recovery	44
Obtaining Automatic System Recovery Information	45
▼ To Retrieve Information About the Status of System Components Affected by ASR	45
Unconfiguring and Reconfiguring Devices	45
▼ To Unconfigure a Device Manually	46
▼ To Reconfigure a Device Manually	47
Displaying System Fault Information	47
▼ To Display Current Valid System Faults	48
▼ To Clear a Fault	48
Storing FRU Information	49
▼ To Store Information in Available FRU PROMs	49

Multipathing Software	49
For More Information	49
3. Managing Disk Volumes	51
OS Patch Requirements	51
Disk Volumes	51
RAID Technology	52
Integrated Stripe Volumes (RAID 0)	52
Integrated Mirror Volumes (RAID 1)	53
Hardware Raid Operations	54
Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks	54
▼ To Create a Hardware Mirrored Volume	55
▼ To Create a Hardware Mirrored Volume of the Default Boot Device	58
▼ To Create a Hardware Striped Volume	60
▼ To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System	61
▼ To Delete a Hardware RAID Volume	63
▼ To Perform a Mirrored Disk Hot-Plug Operation	65
▼ To Perform a Nonmirrored Disk Hot-Plug Operation	66
4. Logical Domains Software	71
About Logical Domains Software	71
Logical Domain Configurations	72
Logical Domains Software Requirements	72
A. Watchdog Timer Application Mode	75
Watchdog Timer Application Mode	75
Watchdog Timer Limitations	76
Using the ntwdt Driver	77

Understanding the User API	78
Using the Watchdog Timer	78
Setting the Timeout Period	78
Enabling or Disabling the Watchdog	79
Rearming the Watchdog	79
Obtaining the State of the Watchdog Timer	80
Finding and Defining Data Structures	80
Example Watchdog Program	80
Watchdog Timer Error Messages	81
B. Alarm Library <code>libtsalarm</code>	83
C. OpenBoot Configuration Variables	87
Index	91

Preface

The *Sun Netra T5440 Server Administration Guide* is for experienced system administrators. The guide includes general descriptive information about the Sun Netra T5440 server from Oracle®, and detailed instructions for configuring and administering the server. To use the information in this document, you must have working knowledge of computer network concepts and terms, and advanced familiarity with Oracle Solaris Operating System.

Note – For information about changing the hardware configuration of your server, or about running diagnostics, see the *Sun Netra T5440 Server Service Manual*.

How This Document Is Organized

- [Chapter 1](#) describes the system console and how to access it.
- [Chapter 2](#) describes the tools used to configure system firmware, including system controller environmental monitoring, Automatic System Recovery (ASR), and multipathing software. In addition, the chapter describes how to unconfigure and reconfigure a device manually.
- [Chapter 3](#) describes redundant array of independent disks (RAID) concepts, and how to configure and manage RAID disk volumes using your server's onboard serial attached SCSI (SAS) disk controller.
- [Chapter 4](#) describes Logical Domain software.
- [Appendix A](#) describes how to configure and use the watchdog timer on the server.
- [Appendix B](#) provides an example program that illustrates how to get or set the status of the alarms.

- [Appendix C](#) provides a list of all OpenBoot configuration variables and a short description of each.

Using UNIX Commands

This document might not contain information on basic UNIX commands and procedures such as shutting down the system, booting the system, and configuring devices. See the following for this information:

- Software documentation that you received with your system
- Solaris OS documentation, which is at:
<http://docs.sun.com>

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#
ILOM service processorz'	->
ALOM compatibility shell	sc>
OpenBoot™ PROM firmwares	ok

Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

Note – Characters display differently depending on browser settings. If characters do not display correctly, change the character encoding in your browser to Unicode UTF-8.

Related Documentation

The following table lists the documentation for this product. The online documentation is available at:

(<http://docs.sun.com/app/docs/prod/server.nebs>)

Application	Title	Part Number	Format	Location
Planning	<i>Sun Netra T5440 Server Site Planning Guide</i>	820-4441	PDF	Online
Installation	<i>Sun Netra T5440 Server Installation Guide</i>	820-4442	PDF, HTML	Online
Administration	<i>Sun Netra T5440 Server Administration Guide</i>	820-4443	PDF, HTML	Online
ILOM Reference	<i>Oracle Integrated Lights Out Manager (ILOM) 2.0 Supplement for the Sun Netra T5440 Server</i>	820-4444	PDF, HTML	Online
ILOM Reference	<i>Oracle Integrated Lights Out Manager (ILOM) 3.0 Supplement for the Sun Netra T5440 Server</i>	820-6891	PDF, HTML	Online
Issues & Updates	<i>Sun Netra T5440 Server Service Manual</i>	820-4445	PDF, HTML	Online
Service	<i>Sun Netra T5440 Server Safety and Compliance Guide</i>	820-4446	PDF, HTML	Online
Compliance	<i>Sun Netra T5440 Server Product Notes</i>	816-4447	PDF, HTML	Online
Overview	<i>Sun Netra T5440 Server Getting Started Guide</i>	820-3016	Printed PDF	Ship kit, Online

Documentation, Support, and Training

TABLE P-1

Sun Function	URL
Documentation	(http://docs.sun.com/)
Support	(http://www.sun.com/support/)
Training	(http://www.sun.com/training/)

Documentation Feedback

Submit comments about this document by clicking the Feedback[+] link at:

<http://docs.sun.com>

Please include the title and part number of your document with your feedback:

Sun Netra T5440 Server Administration Guide, part number 820-4443-11.

Configuring the System Console

This chapter explains what the system console is, describes the different ways of configuring it on Oracle's Sun Netra T5440 server, and helps you understand the relationship between the system console and the service processor. This chapter contains the following sections:

- “Communicating With the System” on page 1
- “Accessing the Service Processor” on page 7
- “Switching Between the Service Processor and the System Console” on page 18
- “ILOM -> Prompt” on page 20
- “OpenBoot ok Prompt” on page 22
- “System Console OpenBoot Configuration Variable Settings” on page 26

Note – For information about changing the hardware configuration of your server, or about running diagnostics, see the *Sun Netra T5440 Server Service Manual*.

Communicating With the System

To install your system software or to diagnose problems, you need some way to interact at a low level with the system. The system console is the facility for doing this. You use the system console to view messages and issue commands. There can be only one system console per computer.

The serial management port (SER MGT) is the default port for accessing the system console upon initial system installation. After installation, you can configure the system console to accept input from and send output to different devices. [TABLE 1-1](#) lists these devices and where they are discussed in this document.

TABLE 1-1 Ways of Communicating With the System

Devices Available	During Installation	After Installation	Further Information
A terminal server attached to the serial management port (SER MGT)	X	X	“Accessing the Service Processor” on page 7
	X	X	“Accessing the System Console Through a Terminal Server” on page 10
	X	X	“System Console OpenBoot Configuration Variable Settings” on page 26
An alphanumeric terminal or similar device attached to the serial management port (SER MGT)	X	X	“Accessing the Service Processor” on page 7
	X	X	“Accessing the System Console Through an Alphanumeric Terminal” on page 15
	X	X	“System Console OpenBoot Configuration Variable Settings” on page 26
A Tip line attached to the serial management port (SER MGT)	X	X	“Accessing the Service Processor” on page 7
	X	X	“Accessing the System Console Through a Tip Connection” on page 12
		X	“Modifying the /etc/remote File” on page 14
	X	X	“System Console OpenBoot Configuration Variable Settings” on page 26
An Ethernet line connected to the network management port (NET MGT)		X	“Activating the Network Management Port” on page 8
A local graphics monitor (graphics accelerator card, graphics monitor, mouse, and keyboard)		X	“Accessing the System Console Through a Local Graphics Monitor” on page 16
		X	“System Console OpenBoot Configuration Variable Settings” on page 26

What the System Console Does

The system console displays status and error messages generated by firmware-based tests during system startup. After those tests run, you can enter special commands that affect the firmware and alter system behavior. For more information about tests that run during the boot process, refer to the *Sun Netra T5440 Server Service Manual* for your server.

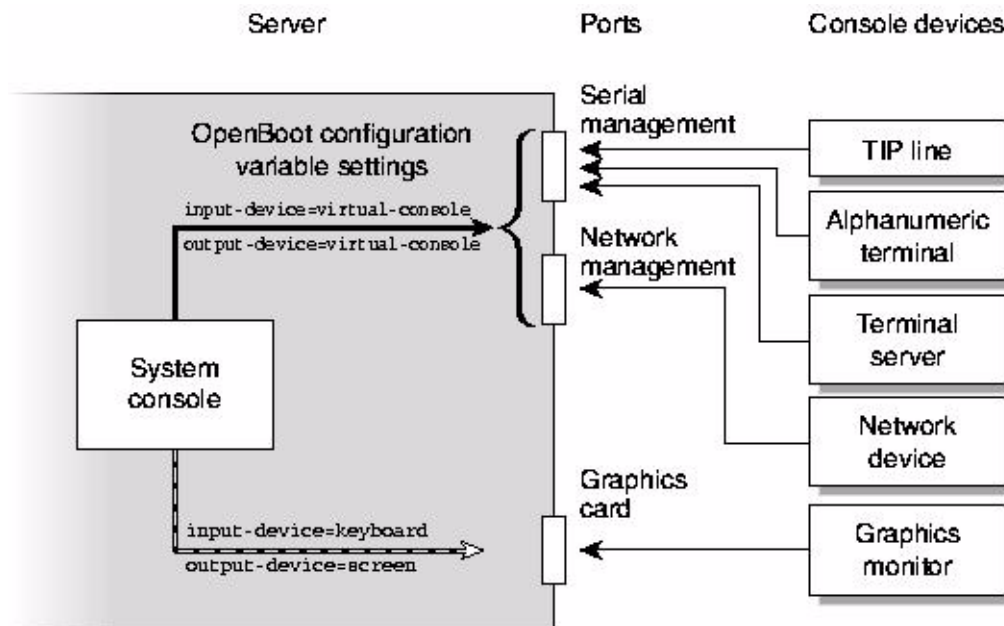
Once the operating system is booted, the system console displays UNIX system messages and accepts UNIX commands.

Using the System Console

To use the system console, you need to attach an input/output device to the system. Initially, you might have to configure that hardware, and load and configure appropriate software as well.

You must also ensure that the system console is directed to the appropriate port on the server's back panel. Generally, the one to which your hardware console device is attached (see [FIGURE 1-1](#)). You do this by setting the `input-device` and `output-device` OpenBoot configuration variables.

FIGURE 1-1 Directing the System Console Input and Output



Default System Console Connection Through the Serial Management and Network Management Ports

On your server, the system console is preconfigured to allow input and output only by means of the service processor. The service processor must be accessed either through the serial management port (SER MGT) or the network management port (NET MGT). By default, the network management port is configured to retrieve network configuration using Dynamic Host Configuration Protocol (DHCP) and to allow connections using Secure Shell (SSH). You can modify the network management port configuration after connecting to ILOM through either the serial or network management ports.

Typically, you connect one of the following hardware devices to the serial management port:

- Terminal server
- Alphanumeric terminal or similar device
- Tip line connected to another computer

These devices provide for secure access at the installation site.

FIGURE 1-2 Rear Panel Connectors, LEDs, and Features on the Sun Netra T5440 Server




Figure Legend

- 1 PCI Slots 0-3: left to right: PCI-X Slot 0 (25 W maximum load), PCI-X Slot 1 (25 W maximum load), PCIe Slot 2 (25 W maximum load), PCIe Slot 3 (25 W maximum load)
- 2 PCI (or XAUI) Slots 4-9: left to right: PCIe or XAUI Slot 4 (15 W maximum load), PCIe or XAUI Slot 5 (15 W maximum load), PCIe Slot 6 (15 W maximum load), PCIe Slot 7 (15 W maximum load), PCIe Slot 8 (15 W maximum load), PCIe Slot 9 (15 W maximum load)
- 3 Service Processor Serial Management Port
- 4 Service Processor Network Management Port
- 5 Gigabit Ethernet Ports left to right: NET0, NET1, NET2, NET3
- 6 Power Supply 0 LEDs top to bottom: Output On LED (green), Service Required LED (yellow), Input Power OK LED (green)
- 7 Power Supplies (PSs): left to right: PS 0, PS 1, PS 2, PS 3
- 8 USB ports left to right: USB2, USB3
- 9 TTYA Serial Port

The service processor serial management port is the default console connection.

Using a Tip line enables you to use windowing and operating system features on the system making the connection to the server.

The serial management port is not a general-purpose serial port. If you want to use a general-purpose serial port with your server, to connect a serial printer for example, use the standard 9-pin serial port on the back panel of the Sun Netra T5440 server. The Solaris OS sees this port as `ttya`.

- For instructions on accessing the system console through a terminal server, see [“Accessing the System Console Through a Terminal Server”](#) on page 10.
- For instructions on accessing the system console through an alphanumeric terminal, see [“Accessing the System Console Through an Alphanumeric Terminal”](#) on page 15.
- For instructions on accessing the system console through a Tip line, see [“Accessing the System Console Through a Tip Connection”](#) on page 12.

Alternative System Console Configuration

In the default configuration, service processor alerts and system console output appear interspersed in the same window. *After initial system installation*, you can redirect the system console to take its input from and send its output to a graphics card's port.

The best practice is to leave the console port in its default configuration for the following reasons:

- In a default configuration, the serial management and network management ports enable you to open up to eight additional windows through which you can view, but not affect, system console activity. You cannot open these connections if the system console is redirected to a graphics card's port.
- In a default configuration, the serial management and network management ports enable you to switch between viewing system console and service processor output on the same device by typing a simple escape sequence or command. The escape sequence and command do not work if the system console is redirected to a graphics card's port.
- The service processor keeps a log of console messages, but some messages are not logged if the system console is redirected to a graphic card's port. The omitted information could be important if you need to contact customer service about a problem.

You change the system console configuration by setting OpenBoot configuration variables. See [“System Console OpenBoot Configuration Variable Settings”](#) on page 26.

Accessing the System Console Through a Graphics Monitor

The Sun Netra T5440 server is shipped without a mouse, keyboard, monitor, or frame buffer for the display of bitmapped graphics. To install a graphics monitor on the server, you must install a graphics accelerator card into a PCI slot, and attach a monitor, mouse, and keyboard to the appropriate front or rear USB ports.

After starting the system, you might need to install the correct software driver for the PCI card you have installed. For detailed hardware instructions, see [“Accessing the System Console Through a Local Graphics Monitor” on page 16](#).

Note – POST diagnostics cannot display status and error messages to a local graphics monitor.

Accessing the Service Processor

The following sections describe methods of accessing the service processor.

Using the Serial Management Port

This procedure assumes that the system console uses the serial management and network management ports (the default configuration).

When you are accessing the system console using a device connected to the serial management port, you first access the ILOM service processor and its `->` prompt. After connecting to the ILOM service processor, you can switch to the system console.

For more information about the ILOM service processor, refer to the ILOM user's guide and the *Integrated Lights Out Management 2.0 Supplement for the Sun Netra T5440 Server*.

▼ To Use the Serial Management Port

1. Ensure that the serial port on your connecting device is set to the following parameters:
 - 9600 baud

- 8 bits
- No parity
- 1 stop bit
- No handshaking

2. Establish an ILOM service processor session.

See the ILOM user's guide for instructions.

3. To connect to the system console, at the ILOM command prompt, type:

```
-> start /SP/console
```

The `start /SP/console` command switches you to the system console.

4. To switch back to the -> prompt, type the #. (Hash-Period) escape sequence.

```
ok #.
```

Characters are not echoed to the screen.

For instructions on how to use the ILOM service processor, refer to the ILOM user's guide and the *Integrated Lights Out Management 2.0 Supplement for the Sun Netra T5440 Server*.

Activating the Network Management Port

The network management port is configured by default to retrieve network settings using DHCP and allow connections using SSH. You might need to modify these settings for your network. If you are unable to use DHCP and SSH on your network, you must connect to the service processor using the serial management port to reconfigure the network management port. See [“Using the Serial Management Port” on page 7](#).

Note – The default username when connecting to the service processor for the first time is `root`. The default password is `changeme`. You should assign a new password during initial system configuration. For more information, refer to your server installation guide, the ILOM user's guide, and the *Integrated Lights Out Management 2.0 Supplement for the Sun Netra T5440 Server*.

You can assign the network management port a static IP address or you can configure the port to obtain an IP address using DHCP from another server.

Data centers frequently devote a separate subnet to system management. If your data center has such a configuration, connect the network management port to this subnet.

Note – The network management port is a 10/100 BASE-T port. The IP address assigned to the network management port is a unique IP address, separate from the main server IP address, and is dedicated for use only with the ILOM service processor.

▼ To Activate the Network Management Port

1. Connect an Ethernet cable to the network management port.
2. Log in to the ILOM service processor through the serial management port.
See the *ILOM 2.0 User's Guide* for instructions.
3. Type one of the following commands:
 - If your network uses static IP addresses, type the following set of commands:

```
-> set /SP/network state=enabled
Set 'state' to 'enabled'

-> set /SP/network pendingipaddress=xx.xxx.xx.xxx
Set 'pendingipaddress' to 'xx.xxx.xx.xxx'

-> set /SP/network pendingipdiscovery=static
Set 'pendingipdiscovery' to 'static'

-> set /SP/network pendingipnetmask=255.255.252.0
Set 'pendingipnetmask' to '255.255.252.0'

-> set /SP/network pendingipgateway=xx.xxx.xx.xxx
Set 'pendingipgateway' to 'xx.xxx.xx.xxx'

-> set /SP/network commitpending=true
Set 'commitpending' to 'true'

->
```

Note – If you have configured your server to use static IP addresses, but you wish to reset your network to use Dynamic Host Configuration Protocol (DHCP), type the following commands:

```
-> set /SP/network pendingipdiscovery=dhcp
Set 'pendingipdiscovery' to 'dhcp'

-> set /SP/network commitpending=true
Set 'commitpending' to 'true'
->
```

4. Type the following command to verify network settings:

```
-> show /SP/network
```

To connect through the network management port, use ssh to the IP address you specified in [Step 3](#).

Accessing the System Console Through a Terminal Server

The following procedure assumes that you are accessing the system console by connecting a terminal server to the serial management port (SER MGT) of your server.

▼ To Access the System Console Through a Terminal Server


- 1. Complete the physical connection from the serial management port to your terminal server.**

The serial management port on the Sun Netra T5440 server is a data terminal equipment (DTE) port. The pinouts for the serial management port correspond with the pinouts for the RJ-45 ports on the serial interface breakout cable supplied by Cisco for use with the Cisco AS2511-RJ terminal server. If you use a terminal server made by another manufacturer, check that the serial port pinouts of the Sun Netra T5440 server matches those of the terminal server you plan to use.

If the pinouts for the server serial ports correspond with the pinouts for the RJ-45 ports on the terminal server, you have two connection options:

- Connect a serial interface breakout cable directly to the Sun Netra T5440 server. See [“Accessing the Service Processor”](#) on page 7.
- Connect a serial interface breakout cable to a patch panel and use the straight-through patch cable (supplied by your server’s manufacturer) to connect the patch panel to the server.

FIGURE 1-3 Patch Panel Connection Between a Terminal Server and a Sun Netra T5440 Server



If the pinouts for the serial management port *do not* correspond with the pinouts for the RJ-45 ports on the terminal server, you need to make a crossover cable that connects each pin on the Sun Netra T5440 server serial management port to the corresponding pin in the terminal server’s serial port.

Serial Port (RJ-45 Connector) Pin	Terminal Server Serial Port Pin
Pin 1 (RTS)	Pin 1 (CTS)
Pin 2 (DTR)	Pin 2 (DSR)
Pin 3 (TXD)	Pin 3 (RXD)
Pin 4 (Signal Ground)	Pin 4 (Signal Ground)
Pin 5 (Signal Ground)	Pin 5 (Signal Ground)
Pin 6 (RXD)	Pin 6 (TXD)
Pin 7 (DSR /DCD)	Pin 7 (DTR)
Pin 8 (CTS)	Pin 8 (RTS)

The following table shows the crossovers that the cable must perform for connecting to a typical terminal server..

2. Open a terminal session on the connecting device, and type:

```
% ssh IP-address-of-terminal-server port-number
```


For example, for a Sun Netra T5440 server connected to port 10000 on a terminal server whose IP address is 192.20.30.10, you would type:

```
% ssh 192.20.30.10 10000
```

Accessing the System Console Through a Tip Connection

Use this procedure to access the Sun Netra T5440 server system console by connecting the serial management port (SER MGT) to the serial port of another system ([FIGURE 1-4](#)).

FIGURE 1-4 Tip Connection Between a Sun Netra T5440 Server and Another System



▼ To Access the System Console Through the Tip Connection

1. Connect the RJ-45 serial cable and, if required, the DB-9 or DB-25 adapter provided.

The cable and adapter connect between another system's serial port (typically `tttyb`) and the serial management port on the back panel of the Sun Netra T5440 server.

2. Ensure that the `/etc/remote` file on the other system contains an entry for `hardwire`.

Most releases of Solaris OS software shipped since 1992 contain an `/etc/remote` file with the appropriate `hardwire` entry. However, if the system is running an older version of Solaris OS software, or if the `/etc/remote` file has been modified, you might need to edit the file. See [“Modifying the `/etc/remote` File” on page 14](#) for details.

3. In a shell tool window on the other system, type:

```
% tip hardwire
```

The system responds by displaying:

```
connected
```

The shell tool is now a Tip window directed to the Sun Netra T5440 server through the system’s serial port. This connection is established and maintained even when the server is completely powered off or just starting up.

Note – Use a shell tool or a terminal (such as `dterm`), not a command tool. Some Tip commands might not work properly in a command tool window.

Modifying the `/etc/remote` File

This procedure might be necessary if you are accessing the Sun Netra T5440 server using a Tip connection from a system running an older version of the Solaris OS software. You might also need to perform this procedure if the `/etc/remote` file on the system has been altered and no longer contains an appropriate `hardwire` entry.

▼ To Modify the `/etc/remote` File

1. Log in as superuser to the system console of a system that you intend to use to establish a Tip connection to your server.
2. Determine the release level of Solaris OS software installed on the system.
Type:

```
# uname -r
```

The system responds with a release number.

3. Take one of the following actions, depending on the number displayed.

- If the number displayed by the `uname -r` command is 5.0 or higher:
- The Solaris OS software shipped with an appropriate entry for hardware in the `/etc/remote` file. If you have reason to suspect that this file was altered and the hardware entry modified or deleted, check the entry against the following example, and edit it as needed.

```
hardware:\
:dv=/dev/term/b:br#9600:el=^C^S^Q^U^D:ie=%$:oe=^D:
```

Note – If you intend to use the system’s serial port A rather than serial port B, edit this entry by replacing `/dev/term/b` with `/dev/term/a`.

- If the number displayed by the `uname -r` command is less than 5.0:
- Check the `/etc/remote` file and add the following entry, if it does not already exist.

```
hardware:\
:dv=/dev/ttyb:br#9600:el=^C^S^Q^U^D:ie=%$:oe=^D:
```

Note – If you intend to use the system’s serial port A rather than serial port B, edit this entry by replacing `/dev/ttyb` with `/dev/ttya`.

The `/etc/remote` file is now properly configured. Continue establishing a Tip connection to the Sun Netra T5440 server system console. See [“Accessing the System Console Through a Tip Connection”](#) on page 12.

If you have redirected the system console to `ttyb` and want to change the system console settings back to use the serial management and network management ports, see [“System Console OpenBoot Configuration Variable Settings”](#) on page 26.

Accessing the System Console Through an Alphanumeric Terminal

Use this procedure when you are accessing the Sun Netra T5440 server system console by connecting the serial port of an alphanumeric terminal to the serial management port (SER MGT) of the server.

▼ To Access the System Console Through an Alphanumeric Terminal

1. **Attach one end of the serial cable to the alphanumeric terminal's serial port.**
Use a null modem serial cable or an RJ-45 serial cable and null modem adapter. Connect this cable to the terminal's serial port connector.
2. **Attach the opposite end of the serial cable to the serial management port on the Sun Netra T5440 server.**
3. **Connect the alphanumeric terminal's power cord to an AC/DC outlet.**
4. **Set the alphanumeric terminal to receive:**
 - 9600 baud
 - 8 bits
 - No parity
 - 1 stop bit
 - No handshake protocol

Refer to the documentation accompanying your terminal for information about how to configure the terminal.

You can now issue system commands and view system messages using the alphanumeric terminal. Continue with your installation or diagnostic procedure, as needed. When you are finished, type the alphanumeric terminal's escape sequence.

For more information about connecting to and using the ILOM service processor, refer to the *ILOM User's Guide* and the *Sun Integrated Lights Out Management (ILOM) x.x Supplement for the Sun Netra T5440 Server*.

Accessing the System Console Through a Local Graphics Monitor

Though it is *not* recommended, the system console can be redirected to the graphics frame buffer. After initial system installation, you can install a local graphics monitor and configure it to access the system console. You *cannot* use a local graphics monitor to perform initial system installation, nor can you use a local graphics monitor to view power-on self-test (POST) messages.

To install a local graphics monitor, you must have the following items:

- Supported PCI-based graphics accelerator card and software driver
- Monitor with appropriate resolution to support the frame buffer
- Supported USB keyboard

- Supported USB mouse

▼ To Access the System Console Through a Local Graphics Monitor

1. Install the graphics card into an appropriate PCI slot.

Installation must be performed by a qualified service provider. For further information, refer to the service manual for your server or contact your qualified service provider.

2. Attach the monitor's video cable to the graphics card's video port.

Tighten the thumbscrews to secure the connection.

3. Connect the monitor's power cord to an AC/DC outlet.

4. Connect the USB keyboard cable to one USB port and the USB mouse cable to the other USB port on the Sun Netra T5440 server back panel (FIGURE 1-2).

5. Obtain the `ok` prompt.

For more information, see [“To Obtain the `ok` Prompt” on page 25](#).

6. Set OpenBoot configuration variables appropriately.

From the existing system console, type:

```
ok setenv input-device keyboard
ok setenv output-device screen
```

Note – There are many other system configuration variables. Although these variables do not affect which hardware device is used to access the system console, some of them affect which diagnostic tests the system runs and which messages the system displays at its console. For details, refer to the *Sun Netra T5440 Server Service Manual*.

7. To cause the changes to take effect, type:

```
ok reset-all
```

The system stores the parameter changes, and boots automatically when the OpenBoot configuration variable `auto-boot?` is set to `true` (the default value).

Note – To cause the parameter changes to take effect, you can also power cycle the system using the front panel Power button.


You can now issue system commands and view system messages using your local graphics monitor. Continue with your installation or diagnostic procedure, as needed.

If you want to redirect the system console back to the serial management and network management ports, see [“System Console OpenBoot Configuration Variable Settings”](#) on page 26.

Switching Between the Service Processor and the System Console

The service processor features two management ports, labeled SER MGT and NET MGT, located on the server’s back panel. If the system console is directed to use the serial management and network management ports (the default configuration), these ports provide access to both the system console and the ILOM command-line interface (the ILOM service processor prompt), each on a separate channel (see [FIGURE 1-5](#)).

FIGURE 1-5 Separate System Console and Service Processor Channels



If the system console is configured to be accessible from the serial management and network management ports, when you connect through one of these ports you can access either the ILOM command-line interface or the system console. You can switch between the ILOM service processor prompt and the system console at any time, but you cannot access both at the same time from a single terminal window or shell tool.

The prompt displayed on the terminal or shell tool tells you which channel you are accessing:

- The # or % prompt indicates that you are at the system console and that the Solaris OS is running.
- The ok prompt indicates that you are at the system console and that the server is running under OpenBoot firmware control.
- The -> prompt indicates that you are at the service processor.

Note – If no text or prompt appears, it might be because no console messages were recently generated by the system. Pressing the terminal's Enter or Return key should produce a prompt.

▼ To Access the System Console Through a Local Graphics Monitor

1. To reach the system console from the service processor, type the `start /SP/console` command at the `->` prompt.
2. To reach the service processor from the system console, type the service processor escape sequence. By default, the escape sequence is `#.` (Pound-Period).

For more information about communicating with the service processor and system console, see:

- [“Communicating With the System” on page 1](#)
- [“ILOM -> Prompt” on page 20](#)
- [“OpenBoot ok Prompt” on page 22](#)
- [“Accessing the Service Processor” on page 7](#)
- *The ILOM User’s Guide and the Sun Integrated Lights Out Management (ILOM) x.x Supplement for the Sun Netra T5440 Server*

ILOM -> Prompt

The ILOM service processor runs independently of the server and regardless of system power state. When you connect your server to AC/DC power, the ILOM service processor immediately starts up, and begins monitoring the system.

Note – To view ILOM service processor boot messages, you must establish a connection using a serial device (such as an Alphanumeric terminal) to the serial management port *before* connecting AC/DC power cords to the server.

You can log in to the ILOM service processor at any time, regardless of system power state, as long as AC/DC power is connected to the system and you have a way of interacting with the system. You can also access the ILOM service processor prompt (`->`) from the OpenBoot `ok` prompt or from the Solaris `#` or `%` prompt, provided the system console is configured to be accessible through the serial management and network management ports.

The `->` prompt indicates that you are interacting with the ILOM service processor directly. The `->` prompt is the first prompt you see when you log in to the system through the serial management port or network management port, regardless of host’s power state.

Note – When you access the ILOM service processor for the first time, the default username is `root` and the default password is `changeme`.

For more information on navigating to or from the ILOM prompt, see the following:

- [“To Obtain the ok Prompt” on page 25](#)
- [“Switching Between the Service Processor and the System Console” on page 18](#)

Access Through Multiple Controller Sessions

Up to five ILOM sessions can be active concurrently, one session through the serial management port and up to four SSH sessions through the network management port. Users of each of these sessions can issue commands at the `->` prompt. However, only one user at a time can access the system console, and then only if the system console is configured to be accessible through the serial and network management ports. For more information, see:

- [“Accessing the Service Processor” on page 7](#)
- [“Activating the Network Management Port” on page 8](#)

Any additional ILOM sessions afford passive views of system console activity, until the active user of the system console logs out.

Reaching the `->` Prompt

There are several ways to obtain the `->` prompt:

- If the system console is directed to the serial management and network management ports, you can type the ILOM escape sequence (`#.`).
- You can log in directly to the service processor from a device connected to the serial management port. See [“Accessing the Service Processor” on page 7](#).
- You can log in directly to the service processor using a connection through the network management port. See [“Activating the Network Management Port” on page 8](#).

OpenBoot ok Prompt

A Sun Netra T5440 server with the Solaris OS installed operates at different run levels. For a full description of run levels, refer to the Solaris system administration documentation.

Most of the time, you operate a Sun Netra T5440 server at run level 2 or run level 3, which are multiuser states with access to full system and network resources. Occasionally, you might operate the system at run level 1, which is a single-user administrative state. However, the lowest operational state is run level 0. At this state, it is safe to turn off power to the system.

When a Sun Netra T5440 server is at run level 0, the `ok` prompt appears. This prompt indicates that the OpenBoot firmware is in control of the system.

There are a number of scenarios under which OpenBoot firmware control can occur.

- By default, before the operating system is installed the system comes up under OpenBoot firmware control.
- When the `auto-boot?` OpenBoot configuration variable is set to `false`, the system boots to the `ok` prompt.
- When the operating system is halted, the system transitions to run level 0 in an orderly way.
- When the operating system crashes, the system reverts to OpenBoot firmware control.
- During the boot process, when there is a serious hardware problem that prevents the operating system from running, the system reverts to OpenBoot firmware control.
- When a serious hardware problem develops while the system is running, the operating system transitions smoothly to run level 0.
- When you deliberately place the system under firmware control in order to execute firmware-based commands, the OpenBoot firmware is in control.

It is the last of these scenarios that most often concerns you as an administrator, since there will be times when you need to reach the `ok` prompt. Several ways to do this are outlined in [“Reaching the ok Prompt” on page 23](#). For detailed instructions, see [“To Obtain the ok Prompt” on page 25](#).

OpenBoot™ ok Prompt Not Available After the Solaris OS Starts

The OpenBoot firmware is not available and might be removed from memory once the Solaris OS starts.

To reach the ok prompt from the Solaris OS, you must first halt the domain. Use the Solaris OS `halt(1M)` command to halt the domain.

Reaching the ok Prompt

There are several ways to reach the ok prompt, depending on the state of the system and the means by which you are accessing the system console. In decreasing order of desirability, these are:

- Graceful shutdown
- ILOM service processor `set /HOST send_break_action=break` and `start /SP/console` command pair
- Break key
- Manual system reset
- Graceful reset of the control domain using ILOM `reset` command

A discussion of each method follows. For step-by-step instructions, see [“To Obtain the ok Prompt” on page 25](#).

Note – As a rule, before suspending the operating system you should back up files, warn users of the impending shutdown, and halt the system in an orderly manner. However, it is not always possible to take such precautions, especially if the system is malfunctioning.

Graceful Shutdown

The preferred method of reaching the ok prompt is to shut down the operating system by issuing an appropriate command (for example, the `shutdown`, `init`, or `uadmin` command) as described in Solaris system administration documentation.

Gracefully shutting down the system prevents data loss, enables you to warn users beforehand, and causes minimal disruption. You can usually perform a graceful shutdown, provided the Solaris OS is running and the hardware has not experienced serious failure.

You can also perform a graceful system shutdown from the ILOM service processor command prompt using the `stop /SYS` command.

▼ To Use ILOM `set /HOST send_break_action=break, start /SP/console` Commands, or Break Key

When it is impossible or impractical to shut down the system gracefully, you can get to the `ok` prompt, if you have an alphanumeric terminal attached to the server, by pressing the Break key.

1. **Type `set /HOST send_break_action=break` to force a running Sun Netra T5440 server to drop to a menu:**

```
-> set /HOST send_break_action=break
Set 'send_break_action' to 'break'
-> start /SP/console
Are you sure you want to start /SP/console (y/n)? y
Serial console started. To stop, type #.
```

2. **Press the Enter key**

The server responds with:

```
c)ontinue, s)ync, r)eboot, h)alt?
```

3. **Type `c` to get OpenBoot firmware control.**

When `init 0` is used, the server drops to this menu:

```
r)eboot, o)k prompt, h)alt?
```

From this menu, you can type `o` to get OpenBoot firmware control.

If the operating system is already halted, you can use the `start /SP/console` command instead of `set /HOST send_break_action=break` to reach the `ok` prompt.

Note – After forcing the system into OpenBoot firmware control, be aware that issuing certain OpenBoot commands (such as `probe-scsi`, `probe-scsi-all`, or `probe-ide`) might hang the system.

Note – These methods of reaching the `ok` prompt will only work if the system console has been redirected to the appropriate port. For details, see [“System Console OpenBoot Configuration Variable Settings”](#) on page 26.

Manual System Reset

Use the ILOM service processor `reset /SYS` command, or `start /SYS` and `stop /SYS` commands, to reset the server. Reaching the `ok` prompt by performing a manual system reset or by power-cycling the system always first attempts a graceful shutdown. If a graceful shutdown is not possible, a forced shutdown is performed. A forced shutdown results in the loss of all system coherence and state information. A forced system reset could corrupt the server's file systems, although the `fsck` command usually restores them.



Caution – Accessing the `ok` prompt suspends the Solaris OS. You cannot return to the Solaris OS without rebooting the OS, for example with the `boot` command.

Graceful Reset of the Control Domain With ILOM reset Command

Use the ILOM `reset` command to gracefully reset the control domain and obtain the `ok` prompt. If a graceful shutdown is not possible, a forced shutdown is performed. To use this method to obtain the `ok` prompt, you must first set the control domain `auto-boot?` option to `false`.

▼ To Obtain the `ok` Prompt

1. Decide which method you need to use to reach the `ok` prompt.

See [“OpenBoot `ok` Prompt” on page 22](#) for details.

2. Follow the appropriate instructions in the following table.

The following table describes ways of accessing the `ok` prompt.

Access Method	What to Do
Graceful shutdown of the Solaris OS	From a shell or command tool window, issue an appropriate command (for example, the <code>shutdown</code> or <code>init</code> command) as described in Solaris system administration documentation.
Break key	From an alphanumeric terminal configured to access the system console, press the Break key.

Access Method	What to Do
ILOM commands	From the -> prompt, type the <code>set /HOST send_break_action=break</code> command. Then issue the <code>start /SP/console</code> command, provided the operating system software is not running and the server is already under OpenBoot firmware control.
Manual system reset	From the -> prompt, type: -> set /HOST/bootmode script="setenv auto-boot? false" Press Enter. then type: -> reset /SYS
Graceful reset of the control domain using the ILOM reset command	From the -> prompt in the control domain, type: -> set /HOST/domain/control auto-boot=disable -> reset /HOST/domain/control

For More Information

For more information about the OpenBoot firmware, refer to the *OpenBoot 4.x Command Reference Manual*. An online version of the manual is available at: <http://docs.sun.com>

System Console OpenBoot Configuration Variable Settings

On the Sun Netra T5440 server, the system console is directed to the serial management and network management ports (SER MGT and NET MGT) by default. However, you can redirect the system console to a local graphics monitor, keyboard, and mouse. You can also redirect the system console back to the serial management and network management ports.

Certain OpenBoot configuration variables control where system console input is taken from and where its output is directed. The table below shows how to set these variables in order to use either the serial management and network management ports, or a local graphics monitor, as the system console connection.

TABLE 1-2 OpenBoot Configuration Variables That Affect the System Console

Setting for Sending System Console Output to:		
OpenBoot Configuration Variable Name	Serial and Network Management Ports	Local Graphics Monitor/USB Keyboard and Mouse
output-device	virtual-console	screen
input-device	virtual-console	keyboard

Note – POST output will still be directed to the serial management port, as POST has no mechanism to direct its output to a graphics monitor.

The serial management port does not function as a standard serial connection. (If you want to connect a conventional serial device (such as a printer) to the system, you must connect it to `ttya` not to the serial management port.)

The `->` prompt and POST messages are only available through the serial management port and network management port. The ILOM service processor `start /SP/console` command is ineffective when the system console is redirected to a local graphics monitor.

In addition to the OpenBoot configuration variables described in [TABLE 1-2](#), there are other variables that affect and determine system behavior. These variables are discussed in more detail in [Appendix A](#).

Managing RAS Features and System Firmware

This chapter describes how to manage reliability, availability, and serviceability (RAS) features and system firmware, including ILOM on the service processor, and automatic system recovery (ASR). In addition, this chapter describes how to unconfigure and reconfigure a device manually, and introduces multipathing software.

This chapter contains the following sections:

- [“Location of the Bezel Server Status and Alarm Status Indicators” on page 34](#)
- [“Status Indicators” on page 31](#)
- [“OpenBoot Emergency Procedures” on page 38](#)
- [“Automatic System Recovery” on page 40](#)
- [“Unconfiguring and Reconfiguring Devices” on page 45](#)
- [“Displaying System Fault Information” on page 47](#)
- [“Multipathing Software” on page 49](#)
- [“Storing FRU Information” on page 49](#)

Note – This chapter does not cover detailed troubleshooting and diagnostic procedures. For information about fault isolation and diagnostic procedures, refer to the *Sun Netra T5440 Server Service Manual*.

ILOM and the Service Processor

The ILOM service processor supports a total of five concurrent sessions per server, four SSH connections available through the network management port and one connection available through the serial management port.

After you log in to your ILOM account, the ILOM service processor command prompt (->) appears, and you can enter ILOM service processor commands. If the command you want to use has multiple options, you can either enter the options individually or grouped together, as shown in the following example.

```
-> stop -force -script /SYS
-> start -script /SYS
```

Logging In To ILOM

All environmental monitoring and control is handled by ILOM on the ILOM service processor. The ILOM service processor command prompt (->) provides you with a way of interacting with ILOM. For more information about the -> prompt, see [“ILOM -> Prompt” on page 20](#).

For instructions on connecting to the ILOM service processor, see:

- [“Accessing the Service Processor” on page 7](#)
- [“Activating the Network Management Port” on page 8](#)

Note – This procedure assumes that the system console is directed to use the serial management and network management ports (the default configuration).

▼ To Log In To ILOM

1. **At the ILOM login prompt, enter the login name and press Return.**

The default login name is root.

```
Integrated Lights Out Manager 2.0
Please login: root
```

2. At the password prompt, enter the password and press Return to get to the -> prompt.

```
Please Enter password:  
->
```

Note – The default user is `root` and the password is `changeme`. For more information, refer to the *Sun Netra T5440 Server Installation Guide*, the *Integrated Lights Out Management User's Guide*, and the *Integrated Lights Out Management 2.0 Supplement for the Sun Netra T5440 Server*.



Caution – To provide optimum system security, change the default system password during initial setup.

Using the ILOM service processor, you can monitor the system, turn the Locator LED on and off, or perform maintenance tasks on the ILOM service processor itself. For more information, refer to the *ILOM User's Guide* and the ILOM supplement for your server.

▼ To View System Fault Information

1. Log in to the ILOM service processor.
2. Use the following command to display a list of faults currently on the system.

```
show /SP/faultmgmt
```

Note – You do not need ILOM administrator permissions to use this command.

Status Indicators

The system has LED indicators associated with the server itself and with various components. The server status indicators are located on the bezel and repeated on the back panel. The components with LED indicators to convey status are the dry contact alarm card, power supply units, Ethernet port, and hard drives.

The topics in this section include:

- [“Interpreting System LEDs” on page 32](#)

- “Bezel Server Status Indicators” on page 33
- “Alarm Status Indicators” on page 35
- “Controlling the Locator LED” on page 37

Interpreting System LEDs

The behavior of LEDs on the Sun Netra T5440 server conform to the American National Standards Institute (ANSI) Status Indicator Standard (SIS). These standard LED behaviors are described in [TABLE 2-1](#).

TABLE 2-1 Standard LED Behaviors and Values

LED Behavior	Meaning
Off	The condition represented by the color is not true.
Steady on	The condition represented by the color is true.
Standby blink	The system is functioning at a minimal level and ready to resume full function.
Slow blink	Transitory activity or new activity represented by the color is taking place.
Fast blink	Attention is required.
Feedback flash	Activity is taking place commensurate with the flash rate (such as disk drive activity).

The system LEDs have assigned meanings, described in [TABLE 2-2](#).

TABLE 2-2 System LED Behaviors With Assigned Meanings

Color	Behavior	Definition	Description
White	Off	Steady state	
	Fast blink	4-Hz repeating sequence, equal intervals On and Off.	This indicator helps you to locate a particular enclosure, board, or subsystem. Example: the Locator LED.
Blue	Off	Steady state	
	Steady on	Steady state	If blue is on, a service action can be performed on the applicable component with no adverse consequences. Example: the OK-to-Remove LED.
Yellow/amber	Off	Steady state	

TABLE 2-2 System LED Behaviors With Assigned Meanings *(Continued) (Continued)*

Color	Behavior	Definition	Description
Green	Slow blink	1-Hz repeating sequence, equal intervals On and Off.	This indicator signals new fault conditions. Service is required. Example: the Service Required LED.
	Steady on	Steady state	The amber indicator stays on until the service action is completed and the system returns to normal function.
	Off	Steady state	
	Standby blink	Repeating sequence consisting of a brief (0.1 sec.) on flash followed by a long off period (2.9 sec.)	The system is running at a minimum level and is ready to be quickly revived to full function. Example: the System Activity LED.
	Steady on	Steady state	Status normal. System or component functioning with no service actions required.
	Slow blink		A transitory (temporary) event is taking place for which direct proportional feedback is not needed or is not feasible.

Bezel Server Status Indicators

[FIGURE 2-1](#) shows the location of the bezel indicators, and [TABLE 2-3](#) provides information about the server status indicators.

FIGURE 2-1 Location of the Bezel Server Status and Alarm Status Indicators




Figure Legend

1	Locator LED and Button	5	User (amber) Alarm Status Indicator
2	Fault LED	6	Minor (amber) Alarm Status Indicator
3	Activity LED	7	Major (red) Alarm Status Indicator
4	Power OK LED	8	Critical (red) Alarm Status Indicator

TABLE 2-3 Bezel Server Status Indicators

Indicator	LED Color	LED State	Component Status
Locator	White	On	Server is identified
		Off	Normal state
Fault	Amber	On	The server has detected a problem and requires the attention of service personnel.

TABLE 2-3 Bezel Server Status Indicators *(Continued)* *(Continued)*

Indicator	LED Color	LED State	Component Status
Activity	Green	Off	The server has no detected faults.
		On	The server is powered up and running the Solaris Operating System.
		Off	Either power is not present or the Solaris software is not running.

Alarm Status Indicators

The dry contact alarm card has four LED status indicators that are supported by ILOM. They are located vertically on the bezel (FIGURE 2-1). Information on the alarm indicators and dry contact alarm states is provided in TABLE 2-4. For more information on alarm indicators, see the *Integrated Lights Out Management User's Guide*.

TABLE 2-4 Alarm Indicators and Dry Contact Alarm States

Indicator and Relay Labels	Indicator Color	Application or Server State	Condition or Action	Activity Indicator State	Alarm Indicator State	Relay NC** State	Relay NO†† State	Comments
Critical (Alarm0)	Red	Server state (Power on or off, and Solaris OS functional or not functional)	No power input	Off	Off	Closed	Open	Default state
			System power off	Off	Off‡	Closed	Open	Input power connected
			System power turns on, Solaris OS not fully loaded	Off	Off‡	Closed	Open	Transient state
			Solaris OS successfully loaded	On	Off	Open	Closed	Normal operating state
			Watchdog timeout	Off	On	Closed	Open	Transient state, reboot Solaris OS
			Solaris OS shutdown initiated by user*	Off	Off‡	Closed	Open	Transient state
			Lost input power	Off	Off	Closed	Open	Default state
			System power shutdown by user	Off	Off‡	Closed	Open	Transient state
		Application state	User sets critical alarm to on†	--	On	Closed	Open	Critical fault detected
			User sets critical alarm to off**	--	Off	Open	Closed	Critical fault cleared
Major (Alarm1)	Red	Application state	User sets major alarm to on**	--	On	Open	Closed	Major fault detected
			User sets major alarm to off‡	--	Off	Closed	Open	Major fault cleared

TABLE 2-4 Alarm Indicators and Dry Contact Alarm States (Continued) (Continued)

Indicator and Relay Labels	Indicator Color	Application or Server State	Condition or Action	Activity Indicator State	Alarm Indicator State	Relay NC** State	Relay NO†† State	Comments
Minor (Alarm2)	Amber	Application state	User sets minor alarm to on**	--	On	Open	Closed	Minor fault detected
			User sets minor alarm to off**	--	Off	Closed	Open	Minor fault cleared
User (Alarm3)	Amber	Application state	User sets user alarm to on**	--	On	Open	Closed	User fault detected
			User sets user alarm to off**	--	Off	Closed	Open	User fault cleared

* The user can shut down the system using commands such as `init0` and `init6`. These commands do not remove power from the system.

† Based on a determination of the fault conditions, the user can turn the alarm on using the Solaris platform alarm API or ILOM CLI.

‡ The implementation of this alarm indicator state is subject to change.

** NC state is the normally closed state. This state represents the default mode of the relay contacts in the normally closed state.

††NO state is the normally open state. This state represents the default mode of the relay contacts in the normally open state.

When the user sets an alarm, a message is displayed on the console. For example, when the critical alarm is set, the following message is displayed on the console:

```
SC Alert: CRITICAL ALARM is set
```

In certain instances when the critical alarm is set, the associated alarm indicator is not lit.

Controlling the Locator LED

You control the Locator LED from the `->` prompt or with the Locator button on the front of the chassis.

▼ To Control the Locator LED

1. To turn on the Locator LED, from the ILOM service processor command prompt, type:

```
-> set /SYS/LOCATE value=on
```

2. To turn off the Locator LED, from the ILOM service processor command prompt, type:

```
-> set /SYS/LOCATE value=off
```

3. To display the state of the Locator LED, from the ILOM service processor command prompt, type:

```
-> show /SYS/LOCATE
```

Note – You do not need Administrator permissions to use the `set /SYS/LOCATE` and `show /SYS/LOCATE` commands

OpenBoot Emergency Procedures

The introduction of Universal Serial Bus (USB) keyboards with the newest systems has made it necessary to change some of the OpenBoot emergency procedures. Specifically, the `Stop-N`, `Stop-D`, and `Stop-F` commands that were available on systems with non-USB keyboards are not supported on systems that use USB keyboards, such as the Sun Netra T5440 server. If you are familiar with the earlier (non-USB) keyboard functionality, this section describes the analogous OpenBoot emergency procedures available in newer systems that use USB keyboards.

OpenBoot Emergency Procedures for the Sun Netra T5440 System

The following sections describe how to perform the functions of the `Stop` commands on systems that use USB keyboards. These same functions are available through Integrated Lights Out Manager (ILOM) system controller software.

Stop-N Functionality

`Stop-N` functionality is not available. However, you can closely emulate the `Stop-N` functionality by completing the following steps, provided the system console is configured to be accessible using either the serial management port or the network management port.

▼ To Restore OpenBoot Configuration Defaults

1. Log in to the ILOM service processor.
2. Type the following commands:

```
-> set /HOST/bootmode state=reset_nvram
-> set /HOST/bootmode script="setenv auto-boot? false"
->
```

Note – If you do not issue the `stop /SYS` and `start /SYS` commands or the `reset /SYS` command within 10 minutes, the host server ignores the `set/HOST/bootmode` commands.

You can issue the `show /HOST/bootmode` command without arguments to display the current setting.

```
-> show /HOST/bootmode

/HOST/bootmode
  Targets:

  Properties:
    config = (none)
    expires = Tue Jan 19 03:14:07 2038
    script = (none)
    state = normal
```

3. To reset the system, type the following commands:

```
-> reset /SYS
Are you sure you want to reset /SYS (y/n)? y
->
```

4. To view console output as the system boots with default OpenBoot configuration variables, switch to console mode.

```
-> set /SP/network pendingipdiscovery=dhcp
Set 'pendingipdiscovery' to 'dhcp'

-> set /SP/network commitpending=true
Set 'commitpending' to 'true'
->
```

5. To discard any customized IDPROM values and restore the default settings for all OpenBoot configuration variables, type:

```
-> set /SP reset_to_defaults=all  
-> reset /SP
```

Stop-F Functionality

The Stop-F functionality is not available on systems with USB keyboards.

Stop-D Functionality

The Stop-D (Diags) key sequence is not supported on systems with USB keyboards. However, you can closely emulate the Stop-D functionality by setting the virtual keyswitch to `diag`, using the ILOM `set /SYS keyswitch_state=diag` command. For more information, refer to the *Integrated Lights Out Management User's Guide* and the *Integrated Lights Out Management 2.0 Supplement for the Sun Netra T5440 Server*.

Automatic System Recovery

The system provides for automatic system recovery (ASR) from failures in memory modules or PCI cards.

Automatic system recovery functionality enables the system to resume operation after experiencing certain nonfatal hardware faults or failures. When ASR is enabled, the system's firmware diagnostics automatically detect failed hardware components. An autoconfiguring capability designed into the system firmware enables the system to unconfigure failed components and to restore system operation. As long as the system is capable of operating without the failed component, the ASR features enable the system to reboot automatically, without operator intervention.

Note – ASR is not activated until you enable it. See [“Enabling and Disabling Automatic System Recovery” on page 43](#).

For more information about ASR, refer to the *Sun Netra T5440 Server Service Manual*.

Auto-Boot Options

The system firmware stores a configuration variable called `auto-boot?`, which controls whether the firmware will automatically boot the operating system after each reset. The default setting for Sun Netra platforms is `true`.

Normally, if a system fails power-on diagnostics, `auto-boot?` is ignored and the system does not boot unless an operator boots the system manually. An automatic boot is generally not acceptable for booting a system in a degraded state. Therefore, the server OpenBoot firmware provides a second setting, `auto-boot-on-error?`. This setting controls whether the system will attempt a degraded boot when a subsystem failure is detected. Both the `auto-boot?` and `auto-boot-on-error?` switches must be set to `true` to enable an automatic degraded boot. To set the switches, type:

```
ok setenv auto-boot? true
ok setenv auto-boot-on-error? true
```

Note – The default setting for `auto-boot-on-error?` is `false`. The system will not attempt a degraded boot unless you change this setting to `true`. In addition, the system will not attempt a degraded boot in response to any fatal nonrecoverable error, even if degraded booting is enabled. For examples of fatal nonrecoverable errors, see [“Error Handling Summary” on page 41](#).

Error Handling Summary

Error handling during the power-on sequence falls into one of the following three cases:

- If no errors are detected by POST or OpenBoot firmware, the system attempts to boot if `auto-boot?` is `true`.
- If only nonfatal errors are detected by POST or OpenBoot firmware, the system attempts to boot if `auto-boot?` is `true` and `auto-boot-on-error?` is `true`. Nonfatal errors include the following:
 - SAS subsystem failure. In this case, a working alternate path to the boot disk is required. For more information, see [“Multipathing Software” on page 49](#).
 - Ethernet interface failure.
 - USB interface failure.
 - Serial interface failure.
 - PCI card failure.

- Memory failure. Given a failed DIMM, the firmware will unconfigure the entire logical bank associated with the failed module. Another nonfailing logical bank must be present in the system for the system to attempt a degraded boot.

Note – If POST or OpenBoot firmware detects a nonfatal error associated with the normal boot device, the OpenBoot firmware automatically unconfigures the failed device and tries the next-in-line boot device, as specified by the `boot-device` configuration variable.

- If a fatal error is detected by POST or OpenBoot firmware, the system does not boot regardless of the settings of `auto-boot?` or `auto-boot-on-error?`. Fatal nonrecoverable errors include the following:
 - Any CPU failed
 - All logical memory banks failed
 - Flash RAM cyclical redundancy check (CRC) failure
 - Critical field-replaceable unit (FRU) PROM configuration data failure
 - Critical system configuration card (SCC) read failure
 - Critical application-specific integrated circuit (ASIC) failure

For more information about troubleshooting fatal errors, refer to the *Sun Netra T5440 Server Service Manual*.

Reset Scenarios

Three ILOM `/HOST/diag` configuration properties, `mode`, `level`, and `trigger`, control whether the system runs firmware diagnostics in response to system reset events.

The standard system reset protocol bypasses POST completely unless the virtual keyswitch or ILOM properties are set as follows:

TABLE 2-5 Virtual Keyswitch Setting for Reset Scenario

Keyswitch	Value
<code>/SYS keyswitch_state</code>	<code>diag</code>

If `keyswitch_state` is set to `diag`, the system can power itself on using preset values of diagnostic properties (`/HOST/diag level=max`, `/HOST/diag mode=max`, `/HOST/diag verbosity=max`) to provide thorough fault coverage. This option overrides the values of diagnostic properties that you might have set elsewhere.

TABLE 2-6 ILOM Property Settings for Reset Scenario

Property	Value
mode	normal or service
level	min or max
trigger	power-on-reset error-reset

The default settings for these properties are:

- `mode = normal`
- `level = min`
- `trigger = power-on-reset error-reset`

For instructions on automatic system recovery (ASR), see [“Enabling and Disabling Automatic System Recovery” on page 43](#).

Automatic System Recovery User Commands

The ILOM commands are available for obtaining ASR status information and for manually unconfiguring or reconfiguring system devices. For more information, see:

- [“Unconfiguring and Reconfiguring Devices” on page 45](#)
- [“To Reconfigure a Device Manually” on page 47](#)
- [“Obtaining Automatic System Recovery Information” on page 45](#)

Enabling and Disabling Automatic System Recovery

The automatic system recovery (ASR) feature is not activated until you enable it. Enabling ASR requires changing configuration variables in ILOM as well as in OpenBoot firmware.

▼ To Enable Automatic System Recovery

1. At the `->` prompt, type:

```
-> set /HOST/diag mode=normal
-> set /HOST/diag level=max
-> set /HOST/diag trigger=power-on-reset
```

2. At the `ok` prompt, type:

```
ok setenv auto-boot true
ok setenv auto-boot-on-error? true
```

Note – For more information about OpenBoot configuration variables, refer to the service manual for your server.

3. To cause the parameter changes to take effect, type:

```
ok reset-all
```

The system permanently stores the parameter changes and boots automatically when the OpenBoot configuration variable `auto-boot?` is set to `true` (its default value).

Note – To store parameter changes, you can also power cycle the system using the front panel Power button.

▼ To Disable Automatic System Recovery

1. At the `ok` prompt, type:

```
ok setenv auto-boot-on-error? false
```

2. To cause the parameter changes to take effect, type:

```
ok reset-all
```

The system permanently stores the parameter change.

Note – To store parameter changes, you can also power cycle the system using the front panel Power button.

After you disable the ASR feature, it is not activated again until you re-enable it.

Obtaining Automatic System Recovery Information

▼ To Retrieve Information About the Status of System Components Affected by ASR

- At the `->` prompt, type:

```
-> show /SYS/component component_state
```

In the `show /SYS/component component_state` command output, any devices marked disabled have been manually unconfigured using the system firmware. The command output also shows devices that have failed firmware diagnostics and have been automatically unconfigured by the system firmware.

Related Information

- [“Automatic System Recovery” on page 40](#)
- [“Enabling and Disabling Automatic System Recovery” on page 43](#)
- [“To Disable Automatic System Recovery” on page 44](#)
- [“Unconfiguring and Reconfiguring Devices” on page 45](#)
- [“To Reconfigure a Device Manually” on page 47](#)

Unconfiguring and Reconfiguring Devices

To support a degraded boot capability, the ILOM firmware provides the `set Device_Identifier component_state=disabled` command, which enables you to unconfigure system devices manually. This command “marks” the specified device as *disabled* by creating an entry in the ASR database. Any device marked

disabled, whether manually or by the system’s firmware diagnostics, is removed from the system’s machine description prior to the hand-off to other layers of system firmware, such as OpenBoot PROM.

▼ To Unconfigure a Device Manually

- At the `->` prompt, type:

`-> set Device-Identifier component_state=disabled`

where the *Device-Identifier* is one of the device identifiers from the following table.

Note – The device identifiers are case sensitive.

Device Identifiers	Devices
<code>/SYS/MB/CMPcpu-number/Pstrand-number</code>	CPU strand (Number: 0-63)
<code>/SYS/MB/PCI_MEZZ/PCIESlot-number</code>	PCIe card (Number: 6-9)
<code>/SYS/MB/PCI_MEZZ/XAUIcard-number</code>	XAUI card (Number: 4-5)
<code>/SYS/MB/PCI_AUX/PCIESlot-number</code>	PCIe card (Number: 0-3)
<code>/SYS/MB/GBEcontroller-number</code>	GBE controllers (Number: 0-1) <ul style="list-style-type: none">• GBE0 controls NET0 and NET1• GBE1 controls NET2 and NET3
<code>/SYS/MB/PCIE</code>	PCIe root complex
<code>/SYS/MB/USBnumber</code>	USB ports (Number: 0-1, located on rear of chassis)
<code>/SYS/MB/CMP0/L2-BANKnumber</code>	(Number: 0-3)
<code>/SYS/DVD</code>	DVD
<code>/SYS/USBBD/USBnumber</code>	USB ports (Number: 2-3, located on front of chassis)
<code>/SYS/TTYA</code>	DB9 serial port
<code>/SYS/MB/CMP0/BRbranch-number/CHchannel-number/Ddimn-number</code>	DIMMS

▼ To Reconfigure a Device Manually

- At the `->` prompt, type:

```
-> set Device-Identifier component-state=enabled
```

where the *Device-Identifier* is any device identifier from the table in the procedure [“To Unconfigure a Device Manually”](#) on page 46.

Note – The device identifiers are not case sensitive. You can type them as uppercase or lowercase characters.

You can use the ILOM `set Device-Identifier component_state=enabled` command to reconfigure any device that you previously unconfigured with the `set Device-Identifier component_state=disabled` command.

Displaying System Fault Information

ILOM software enables you to display current valid system faults.

▼ To Display Current Valid System Faults

- Type:

```
-> show /SP/faultmgmt
```

This command displays the fault ID, the faulted FRU device, and the fault message to standard output. The `show /SP/faultmgmt` command also displays POST results.

For example:

```
-> show /SP/faultmgmt  
  /SP/faultmgmt  
    Targets:  
      0 (/SYS/PS1)  
  
    Properties:  
  
    Commands:  
      cd  
      show  
->
```

For more information about the `show /SP/faultmgmt` command, refer to the ILOM guide and the ILOM supplement for your server.

▼ To Clear a Fault

- Type:

```
-> set /SYS/component clear_fault_action=true
```

Setting `clear_fault_action` to `true` clears the fault at the component and all levels below it in the `/SYS` tree.

Storing FRU Information

▼ To Store Information in Available FRU PROMs

- At the `->` prompt type:

```
-> set /SP customer_frudata=data
```

Multipathing Software

Multipathing software enables you to define and control redundant physical paths to I/O devices such as storage devices and network interfaces. If the active path to a device becomes unavailable, the software can automatically switch to an alternate path to maintain availability. This capability is known as *automatic failover*. To take advantage of multipathing capabilities, you must configure the server with redundant hardware, such as redundant network interfaces or two host bus adapters connected to the same dual-ported storage array.

For the Sun Netra T5440 Server, three different types of multipathing software are available:

- Solaris IP Network Multipathing software provides multipathing and load-balancing capabilities for IP network interfaces.
- VERITAS Volume Manager (VVM) software includes a feature called Dynamic Multipathing (DMP), which provides disk multipathing as well as disk load balancing to optimize I/O throughput.
- Sun StorageTek™ Traffic Manager is an architecture fully integrated within the Solaris OS (beginning with the Solaris 8 release) that enables I/O devices to be accessed through multiple host controller interfaces from a single instance of the I/O device.

For More Information

For instructions on how to configure and administer Solaris IP Network Multipathing, consult the *IP Network Multipathing Administration Guide* provided with your specific Solaris release.

For information about VVM and its DMP feature, refer to the documentation provided with the VERITAS Volume Manager software.

For information about Sun StorageTek Traffic Manager, refer to your Solaris OS documentation.

Managing Disk Volumes

This chapter describes redundant array of independent disks (RAID) concepts, and how to configure and manage RAID disk volumes using the Sun Netra T5440 server's on-board serial attached SCSI (SAS) disk controller.

This chapter contains the following sections:

- [“OS Patch Requirements” on page 51](#)
- [“Disk Volumes” on page 51](#)
- [“RAID Technology” on page 52](#)
- [“To Perform a Nonmirrored Disk Hot-Plug Operation” on page 66](#)

OS Patch Requirements

To configure and use RAID disk volumes on the Sun Netra T5440 server, you must install the appropriate patches. For the latest information on patches for the Sun Netra T5440 server, see the latest product notes for your system.

Installation procedures for patches are included in text README files that accompany the patches.

Disk Volumes

From the perspective of the on-board disk controller on the Sun Netra T5440 server, *disk volumes* are logical disk devices comprising one or more complete physical disks.

Once you create a volume, the operating system uses and maintains the volume as if it were a single disk. By providing this logical volume management layer, the software overcomes the restrictions imposed by physical disk devices.

The onboard disk controller of the Sun Netra T5440 server provides for the creation of as many as two hardware RAID volumes. The controller supports either two-disk RAID 1 (integrated mirror, or IM) volumes, or up to eight-disk RAID 0 (integrated stripe, or IS) volumes.

Note – Due to the volume initialization that occurs on the disk controller when a new volume is created, properties of the volume such as geometry and size are unknown. RAID volumes created using the hardware controller must be configured and labeled using `format(1M)` prior to use with the Solaris Operating System. See [“To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System” on page 61](#), or the `format(1M)` man page for further details.

Volume migration (relocating all RAID volume disk members from one Sun Netra T5440 chassis to another) is not supported. If you must perform this operation, contact your service provider.

RAID Technology

RAID technology enables the construction of a logical volume, made up of several physical disks, in order to provide data redundancy, increased performance, or both. The Sun Netra T5440 server’s on-board disk controller supports both RAID 0 and RAID 1 volumes.

This section describes the RAID configurations supported by the on-board disk controller:


- Integrated Stripe, or IS volumes (RAID 0)
- Integrated Mirror, or IM volumes (RAID 1)

Integrated Stripe Volumes (RAID 0)

Integrated Stripe volumes are configured by initializing the volume across two or more physical disks, and sharing the data written to the volume across each physical disk in turn, or *striping* the data across the disks.

Integrated Stripe volumes provide for a logical unit (LUN) that is equal in capacity to the sum of all its member disks. For example, a three-disk IS volume configured on 72-gigabyte drives will have a capacity of 216 gigabytes.

FIGURE 3-1 Graphical Representation of Disk Striping




Caution – There is no data redundancy in an IS volume configuration. Thus, if a single disk fails, the entire volume fails, and all data is lost. If an IS volume is manually deleted, all data on the volume is lost.

IS volumes are likely to provide better performance than IM volumes or single disks. Under certain workloads, particularly some write or mixed read-write workloads, I/O operations complete faster because the I/O operations are being handled in a round-robin fashion, with each sequential block being written to each member disk in turn.

Integrated Mirror Volumes (RAID 1)

Disk mirroring (RAID 1) is a technique that uses data redundancy, two complete copies of all data stored on two separate disks, to protect against loss of data due to disk failure. One logical volume is duplicated on two separate disks.

FIGURE 3-2 Graphical Representation of Disk Mirroring



Whenever the operating system needs to write to a mirrored volume, both disks are updated. The disks are maintained at all times with exactly the same information. When the operating system needs to read from the mirrored volume, the OS reads from whichever disk is more readily accessible at the moment. This functionality can result in enhanced performance for read operations.



Caution – Creating RAID volumes using the on-board disk controller destroys all data on the member disks. The disk controller’s volume initialization procedure reserves a portion of each physical disk for metadata and other internal information used by the controller. Once the volume initialization is complete, you can configure the volume and label it using the `format(1M)` utility. You can then use the volume in the Solaris OS.

Hardware Raid Operations

On the Sun Netra T5440 server, the SAS controller supports mirroring and striping using the Solaris OS `raidctl` utility.

A hardware RAID volume created under the `raidctl` utility behaves slightly differently than one created using volume management software. Under a software volume, each device has its own entry in the virtual device tree, and read-write operations are performed to both virtual devices. Under hardware RAID volumes, only one device appears in the device tree. Member disk devices are invisible to the operating system, and are accessed only by the SAS controller.

Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks

To perform a disk hot-plug procedure, you must know the physical or logical device name for the drive that you want to install or remove. If your system encounters a disk error, often you can find messages about failing or failed disks in the system console. This information is also logged in the `/var/adm/messages` files.

These error messages typically refer to a failed hard drive by its physical device name (such as `/devices/pci@1f,700000/scsi@2/sd@1,0`) or by its logical device name (such as `c0t1d0`). In addition, some applications might report a disk slot number (0 through 3).

You can use [TABLE 3-1](#) to associate internal disk slot numbers with the logical and physical device names for each hard drive.

TABLE 3-1 Disk Slot Numbers, Logical Device Names, and Physical Device Names

Disk Slot Number	Logical Device Name*	Physical Device Name
Slot 0	c0t0d0	/devices/pci@0/pci@0/pci@2/scsi@0/sd@0,0
Slot 1	c0t1d0	/devices/pci@0/pci@0/pci@2/scsi@0/sd@1,0
Slot 2	c0t2d0	/devices/pci@0/pci@0/pci@2/scsi@0/sd@2,0
Slot 3	c0t3d0	/devices/pci@0/pci@0/pci@2/scsi@0/sd@3,0

* The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Create a Hardware Mirrored Volume

1. Verify which hard drive corresponds with which logical device name and physical device name, using the `raidctl` command:

```
# raidctl
No RAID volumes found.
```

See [“Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks”](#) on page 54.

The preceding example indicates that no RAID volume exists. In another case:

```
# raidctl
RAID      Volume  RAID      RAID      Disk
Volume    Type    Status    Disk      Status
-----
c0t0d0    IM      OK              c0t0d0      OK
                               c0t1d0      OK
```

In this example, a single IM volume has been enabled. It is fully synchronized and is online.

The Sun Netra T5440 server’s on-board SAS controller can configure as many as two RAID volumes. Prior to volume creation, ensure that the member disks are available and that there are not two volumes already created.

The RAID status might be:

- OK – Indicating that the RAID volume is online and fully synchronized.
- RESYNCING – Indicating that the data between the primary and secondary member disks in an IM are still synchronizing.

- **DEGRADED** – Indicating that a member disk is failed or otherwise offline.
- **FAILED** – Indicating that volume should be deleted and reinitialized. This failure can occur when any member disk in an IS volume is lost, or when both disks are lost in an IM volume.

The Disk Status column displays the status of each physical disk. Each member disk might be **OK**, indicating that it is online and functioning properly, or it might be **FAILED**, **MISSING**, or otherwise **OFFLINE**, indicating that the disk has hardware or configuration issues that need to be addressed.

For example, an IM with a secondary disk that has been removed from the chassis appears as:

# raidctl				
RAID	Volume	RAID	RAID	Disk
Volume	Type	Status	Disk	Status

c0t0d0	IM	DEGRADED	c0t0d0	OK
			c0t1d0	MISSING

See the `raidctl(1M)` man page for additional details regarding volume and disk status.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

2. Type the following command:

```
# raidctl -c primary secondary
```

The creation of the RAID volume is interactive, by default. For example:

```
# raidctl -c c0t0d0 c0t1d0  
Creating RAID volume c0t0d0 will destroy all data on member disks,  
proceed  
(yes/no)? yes  
Volume 'c0t0d0' created  
#
```

As an alternative, you can use the `-f` option to force the creation if you are sure of the member disks, and sure that the data on both member disks can be lost. For example:

```
# raidctl -f -c c0t0d0 c0t1d0  
Volume 'c0t0d0' created  
#
```

When you create a RAID mirror, the secondary drive (in this case, `c0t1d0`) disappears from the Solaris device tree.

3. To check the status of a RAID mirror, type the following command:

# raidctl				
RAID	Volume	RAID	RAID	Disk
Volume	Type	Status	Disk	Status

c0t0d0	IM	RESYNCING	c0t0d0	OK
			c0t1d0	OK

The preceding example indicates that the RAID mirror is still resynchronizing with the backup drive.

The following example shows that the RAID mirror is synchronized and online.

# raidctl				
RAID	Volume	RAID	RAID	Disk
Volume	Type	Status	Disk	Status

c0t0d0	IM	OK	c0t0d0	OK
			c0t1d0	OK

The disk controller synchronizes IM volumes one at a time. If you create a second IM volume before the first IM volume completes its synchronization, the first volume's RAID status will indicate RESYNCING, and the second volume's RAID status will indicate OK. Once the first volume has completed, its RAID status changes to OK, and the second volume automatically starts synchronizing, with a RAID status of RESYNCING.

Under RAID 1 (disk mirroring), all data is duplicated on both drives. If a disk fails, replace it with a working drive and restore the mirror. For instructions, see ["To Perform a Mirrored Disk Hot-Plug Operation" on page 65](#).

For more information about the `raidctl` utility, see the `raidctl(1M)` man page.

▼ To Create a Hardware Mirrored Volume of the Default Boot Device

Due to the volume initialization that occurs on the disk controller when a new volume is created, the volume must be configured and labeled using the `format(1M)` utility prior to use with the Solaris Operating System (see ["To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System" on page 61](#)). Because of this limitation, `raidctl(1M)` blocks the creation of a hardware RAID volume if any of the member disks currently have a file system mounted.

This section describes the procedure required to create a hardware RAID volume containing the default boot device. Since the boot device always has a mounted file system when booted, an alternate boot medium must be employed, and the volume

created in that environment. One alternate medium is a network installation image in single-user mode (refer to the *Solaris 10 Installation Guide* for information about configuring and using network-based installations).

1. Determine which disk is the default boot device

From the OpenBoot ok prompt, type the `printenv` command, and if necessary the `devalias` command, to identify the default boot device. For example:

```
ok printenv boot-device
boot-device =          disk

ok devalias disk
disk                  /pci@0/pci@0/pci@2/scsi@0/disk@0,0
```

2. Type the `boot net -s` command

```
ok boot net -s
```

3. Once the system has booted, use the `raidctl(1M)` utility to create a hardware mirrored volume, using the default boot device as the primary disk.

See [“To Create a Hardware Mirrored Volume” on page 55](#). For example:

```
# raidctl -c -r 1 c0t0d0 c0t1d0
Creating RAID volume c0t0d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume  c0t0d0  created
#
```

4. Install the volume with the Solaris OS using any supported method.

The hardware RAID volume `c0t0d0` appears as a disk to the Solaris installation program.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Create a Hardware Striped Volume

1. Verify which hard drive corresponds with which logical device name and physical device name.

See “Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks” on page 54.

To verify the current RAID configuration, type:

```
# raidctl
No RAID volumes found.
```

The preceding example indicates that no RAID volume exists.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

2. Type the following command:

```
# raidctl -c -r 0 disk1 disk2 ...
```

The creation of the RAID volume is interactive, by default. For example:

```
# raidctl -c -r 0 c0t1d0 c0t2d0 c0t3d0
Creating RAID volume c0t1d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume 'c0t1d0' created
#
```

When you create a RAID striped volume, the other member drives (in this case, c0t2d0 and c0t3d0) disappear from the Solaris device tree.

As an alternative, you can use the `-f` option to force the creation if you are sure of the member disks, and sure that the data on all other member disks can be lost. For example:

```
# raidctl -f -c -r 0 c0t1d0 c0t2d0 c0t3d0
Volume 'c0t1d0' created
#
```

3. To check the status of a RAID striped volume, type the following command:

```
# raidctl
RAID      Volume  RAID      RAID      Disk
Volume   Type    Status    Disk      Status
-----
c0t1d0   IS      OK              c0t1d0     OK
              c0t2d0     OK
              c0t3d0     OK
```

The example shows that the RAID striped volume is online and functioning.

Under RAID 0 (disk striping), there is no replication of data across drives. The data is written to the RAID volume across all member disks in a round-robin fashion. If any one disk is lost, all data on the volume is lost. For this reason, RAID 0 cannot be used to ensure data integrity or availability, but can be used to increase write performance in some scenarios.

For more information about the `raidctl` utility, see the `raidctl(1M)` man page.

▼ To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System

After a creating a RAID volume using `raidctl`, use `format(1M)` to configure and label the volume before attempting to use it in the Solaris OS.

1. Start the `format` utility:

```
# format
```

The `format` utility might generate messages about corruption of the current label on the volume, which you are going to change. You can safely ignore these messages.

2. Select the disk name that represents the RAID volume that you have configured.

In this example, `c0t2d0` is the logical name of the volume.

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
    0. c0t0d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
       /pci@0/pci@0/pci@2/scsi@0/sd@0,0
    1. c0t1d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
       /pci@0/pci@0/pci@2/scsi@0/sd@1,0
    2. c0t2d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
       /pci@0/pci@0/pci@2/scsi@0/sd@2,0
```

```
Specify disk (enter its number): 2
selecting c0t2d0
[disk formatted]
FORMAT MENU:
    disk          - select a disk
    type          - select (define) a disk type
    partition     - select (define) a partition table
    current       - describe the current disk
    format        - format and analyze the disk
    fdisk         - run the fdisk program
    repair        - repair a defective sector
    label         - write label to the disk
    analyze       - surface analysis
    defect        - defect list management
    backup        - search for backup labels
    verify        - read and display labels
    save          - save new disk/partition definitions
    inquiry       - show vendor, product and revision
    volname       - set 8-character volume name
    !<cmd>        - execute <cmd>, then return
quit
```

3. Type the `type` command at the `format>` prompt, then select 0 (zero) to auto configure the volume.

For example:

```
format> type

AVAILABLE DRIVE TYPES:
    0. Auto configure
    1. DEFAULT
    2. SUN72G
    3. SUN72G
    4. other
Specify disk type (enter its number)[3]: 0
c0t2d0: configured with capacity of 68.23GB
<LSILOGIC-LogicalVolume-3000 cyl 69866 alt 2 hd 16 sec 128>
selecting c0t2d0
[disk formatted]
```

4. Use the `partition` command to partition, or *slice*, the volume according to your desired configuration.

See the `format(1M)` man page for additional details.

5. Write the new label to the disk using the `label` command.

```
format> label  
Ready to label disk, continue? yes
```

6. Verify that the new label has been written by printing the disk list using the `disk` command.

```
format> disk  
  
AVAILABLE DISK SELECTIONS:  
    0. c0t0d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>  
       /pci@0/pci@0/pci@2/scsi@0/sd@0,0  
    1. c0t1d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>  
       /pci@0/pci@0/pci@2/scsi@0/sd@1,0  
    2. c0t2d0 <LSILOGIC-LogicalVolume-3000 cyl 69866 alt 2 hd  
16 sec 128>  
       /pci@0/pci@0/pci@2/scsi@0/sd@2,0  
Specify disk (enter its number) [2]:
```

Note that `c0t2d0` now has a type indicating it is an `LSILOGIC-LogicalVolume`.

7. Exit the `format` utility.

The volume can now be used in the Solaris OS.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Delete a Hardware RAID Volume

1. Verify which hard drive corresponds with which logical device name and physical device name.

See “Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks” on page 54.

2. Determine the name of the RAID volume, type:

```
# raidctl
RAID      Volume  RAID      RAID      Disk
Volume   Type    Status    Disk      Status
-----
c0t0d0   IM      OK        c0t0d0    OK
                   c0t1d0    OK
```

In this example, the RAID volume is c0t1d0.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

3. To delete the volume, type the following command:

```
# raidctl -d mirrored-volume
```

For example:

```
# raidctl -d c0t0d0
RAID Volume 'c0t0d0' deleted
```

In the event that the RAID volume is an IS volume, the deletion of the RAID volume is interactive, for example:

```
# raidctl -d c0t0d0
Deleting volume c0t0d0 will destroy all data it contains, proceed
(yes/no)? yes
Volume 'c0t0d0' deleted.
#
```

The deletion of an IS volume results in the loss of all data that it contains. As an alternative, you can use the `-f` option to force the deletion if you are sure that you no longer need the IS volume, or the data it contains. For example:

```
# raidctl -f -d c0t0d0
Volume 'c0t0d0' deleted.
#
```

4. To confirm that you have deleted the RAID array, type the following command:

```
# raidctl
```

For example:

```
# raidctl
No RAID volumes found
```

For more information, see the `raidctl(1M)` man page.

▼ To Perform a Mirrored Disk Hot-Plug Operation

1. Verify which hard drive corresponds with which logical device name and physical device name.

See “Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks” on page 54.

2. To confirm a failed disk, type the following command:

```
# raidctl
```

If the Disk Status is FAILED, then the drive can be removed and a new drive inserted. Upon insertion, the new disk should be OK and the volume should be RESYNCING.

For example:

# raidctl				
RAID	Volume	RAID	RAID	Disk
Volume	Type	Status	Disk	Status

c0t1d0	IM	DEGRADED	c0t1d0	OK
			c0t2d0	FAILED

This example indicates that the disk mirror has degraded due to a failure in disk `c0t2d0`.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

3. Remove the hard drive, as described in your *Sun Netra T5440 Server Service Manual*.

There is no need to issue a software command to bring the drive offline when the drive has failed.

4. Install a new hard drive, as described in your *Sun Netra T5440 Server Service Manual*.

The RAID utility automatically restores the data to the disk.

5. To check the status of a RAID rebuild, type the following command:

```
# raidctl
```

For example:

```
# raidctl
RAID      Volume  RAID      RAID      Disk
Volume    Type    Status    Disk      Status
-----
c0t1d0    IM      RESYNCING  c0t1d0    OK
                               c0t2d0    OK
```

This example indicates that RAID volume c0t1d0 is resynchronizing.

If you issue the command again once synchronization has completed, it indicates that the RAID mirror is finished resynchronizing and is back online:

```
# raidctl
RAID      Volume  RAID      RAID      Disk
Volume    Type    Status    Disk      Status
-----
c0t1d0    IM      OK        c0t1d0    OK
                               c0t2d0    OK
```

For more information, see the *raidctl(1M)* man page.

▼ To Perform a Nonmirrored Disk Hot-Plug Operation

1. Verify which hard drive corresponds with which logical device name and physical device name.

See [“Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks”](#) on page 54.

Ensure that no applications or processes are accessing the hard drive.

2. Type the following command:

```
# cfmadm -al
```

For example:

```
# cfmadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             scsi-bus      connected   configured  unknown
c0::disk/c0t0d0 disk         connected   configured  unknown
c0::disk/c0t1d0 disk         connected   configured  unknown
c0::disk/c0t2d0 disk         connected   configured  unknown
c0::disk/c0t3d0 disk         connected   configured  unknown
c1             scsi-bus      connected   configured  unknown
c1::disk/c1t0d0 CD-ROM       connected   configured  unknown
usb0/1         unknown      empty       unconfigured ok
usb0/2         unknown      empty       unconfigured ok
usb1/1.1       unknown      empty       unconfigured ok
usb1/1.2       unknown      empty       unconfigured ok
usb1/1.3       unknown      empty       unconfigured ok
usb1/1.4       unknown      empty       unconfigured ok
usb1/2         unknown      empty       unconfigured ok
#
```

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

The -al options return the status of all SCSI devices, including buses and USB devices. In this example, no USB devices are connected to the system.

Note – While you can use the Solaris OS `cfgadm install_device` and `cfgadm remove_device` commands to perform a hard drive hot-plug procedure, these commands issue the following warning message when you invoke them on a bus containing the system disk:

```
# cfgadm -x remove_device c0::dsk/c0t1d0
Removing SCSI device: /devices/pci@1f,4000/scsi@3/sd@1,0
This operation will suspend activity on SCSI bus: c0
Continue (yes/no)? y
dev = /devices/pci@780/pci@0/pci@9/scsi@0/sd@1,0
cfgadm: Hardware specific failure: failed to suspend:
      Resource              Information
-----
/dev/dsk/c0t0d0s0    mounted filesystem "/"
/dev/dsk/c0t0d0s6    mounted filesystem "/usr"
```

This warning is issued because these commands attempt to quiesce the (SAS) SCSI bus, but the server firmware prevents it. This warning message can be safely ignored in the Sun Netra T5440 server, but the following step avoids this warning message altogether.

3. Remove the hard drive from the device tree.

To remove the hard drive from the device tree, type the following command:

```
# cfgadm -c unconfigure Ap-ld
```

For example:

```
# cfgadm -c unconfigure c0::dsk/c0t3d0
```

This example removes `c0t3d0` from the device tree. The blue OK-to-Remove LED lights.

4. Verify that the device has been removed from the device tree.

Type the following command:

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             scsi-bus      connected   configured  unknown
c0::dsk/c0t0d0 disk         connected   configured  unknown
c0::dsk/c0t1d0 disk         connected   configured  unknown
c0::dsk/c0t2d0 disk         connected   configured  unknown
c0::dsk/c0t3d0 unavailable  connected   configured  unknown
c1             scsi-bus      connected   unconfigured unknown
c1::dsk/c1t0d0 CD-ROM       connected   configured  unknown
usb0/1         unknown      empty       unconfigured ok
usb0/2         unknown      empty       unconfigured ok
usb1/1.1       unknown      empty       unconfigured ok
usb1/1.2       unknown      empty       unconfigured ok
usb1/1.3       unknown      empty       unconfigured ok
usb1/1.4       unknown      empty       unconfigured ok
usb1/2         unknown      empty       unconfigured ok
#
```

Note that c0t3d0 is now unavailable and unconfigured. The corresponding hard drive OK-to-Remove LED is lit.

5. Remove the hard drive, as described in *Sun Netra T5440 Server Service Manual*.

The blue OK-to-Remove LED goes out when you remove the hard drive.

6. Install a new hard drive, as described in *Sun Netra T5440 Server Service Manual*.

7. Configure the new hard drive.

Type the following command:

```
# cfgadm -c configure Ap-Id
```

For example:

```
# cfgadm -c configure c1::dsk/c0t3d0
```

The green Activity LED flashes as the new disk at c1t3d0 is added to the device tree.

8. Verify that the new hard drive is in the device tree.

Type the following command:

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             scsi-bus      connected   configured  unknown
```

c0::dsk/c0t0d0	disk	connected	configured	unknown
c0::dsk/c0t1d0	disk	connected	configured	unknown
c0::dsk/c0t2d0	disk	connected	configured	unknown
c0::dsk/c0t3d0	disk	connected	configured	unknown
c1	scsi-bus	connected	configured	unknown
c1::dsk/c1t0d0	CD-ROM	connected	configured	unknown
usb0/1	unknown	empty	unconfigured	ok
usb0/2	unknown	empty	unconfigured	ok
usb1/1.1	unknown	empty	unconfigured	ok
usb1/1.2	unknown	empty	unconfigured	ok
usb1/1.3	unknown	empty	unconfigured	ok
usb1/1.4	unknown	empty	unconfigured	ok
usb1/2	unknown	empty	unconfigured	ok
#				

Note – c0t3d0 is now listed as configured.

Logical Domains Software

Oracle's Sun Netra T5440 server supports the Logical Domains (LDoms) 1.0.3. software that is used to create and manage logical domains. The software comprises LDoms-enabling code in the Solaris 10 5/08 OS, and the Logical Domains Manager, which is the command-line interface.

This chapter includes the following topics:

- [“About Logical Domains Software” on page 71](#)
- [“Logical Domain Configurations” on page 72](#)
- [“Logical Domains Software Requirements” on page 72](#)

About Logical Domains Software

Logical Domains software enables you to allocate the system resources of your server (such as a boot environment, CPUs, memory, and I/O devices) into logical domains. By using a logical domains environment, you can increase resource usage, improve scaling, and gain greater control of security and isolation.

LDoms software enables you to create and manage as many as 128 logical domains, depending on the hardware configuration of the server on which the Logical Domains Manager has been installed. You can virtualize resources and define network, storage, and other I/O devices as services that can be shared between domains.

A logical domain is a discrete logical grouping with its own operating systems, resources, and identity within a single computer system. Applications software can run in logical domains. Each logical domain can be created, destroyed, reconfigured, and rebooted independently, without requiring a power cycle of the server. There are several roles that logical domains can perform as shown in the following table.

TABLE 4-1 Logical Domain Roles

Domain Role	Description
Control domain	Domain in which the Logical Domains Manager runs, enabling you to create and manage other logical domains and allocate virtual resources to other domains. There can be only one control domain per server. The initial domain created when installing Logical Domains software is a control domain and is named <code>primary</code> .
Service domain	Domain that provides virtual device services to other domains, such as a virtual switch, a virtual console concentrator, and a virtual disk server.
I/O domain	Domain that has direct ownership of and direct access to physical I/O devices, such as a network card in a PCI Express controller. Shares the devices to other domains in the form of virtual devices. You can have a maximum of two I/O domains, one of which also must be the control domain.
Guest domain	Domain that is managed by the control domain and uses services from the I/O and service domains.

Logical Domain Configurations

The current configuration of a logical domain can be stored on the service processor (SP). Using Logical Domains Manager CLI commands, you can add a configuration, specify a configuration to be used, and list the configurations on the service processor. You can also use the ILOM `set /HOST/bootmode config=configfile` command to specify an LDom's boot configuration. For further information about `/HOST/bootmode`, refer to the *Integrated Lights Out Management 2.0 Supplement for the Sun Netra T5440 Server*.

Logical Domains Software Requirements

The following software is required or suggested for Logical Domains on the Sun Netra T5440 server:

- (Required) Logical Domains Manager 1.0.3 or above software
- (Suggested) Solaris Security Toolkit 4.2 software – Refer to the *Solaris Security Toolkit 4.2 Administration Guide* and *Solaris Security Toolkit 4.2 Reference Manual* for more information

Note – During the boot process, domains that use virtual boot devices must wait for their service domains to come online first. This can prolong the boot process.

Watchdog Timer Application Mode

This appendix gives information on the watchdog timer application mode on the server. It provides the following sections to help you understand how to configure and use the watchdog timer:

- “Watchdog Timer Application Mode” on page 75
- “Watchdog Timer Limitations” on page 76
- “Using the `ntwdt` Driver” on page 77
- “Understanding the User API” on page 78
- “Using the Watchdog Timer” on page 78
- “Watchdog Timer Error Messages” on page 81

Note – Once the application watchdog timer is in use, you must reboot the Solaris OS in order to return to the default (nonprogrammable) watchdog timer and default LED behavior (no Alarm3).

Watchdog Timer Application Mode

The watchdog mechanism detects a system hang, or an application hang or crash, should they occur. The watchdog is a timer that is continually reset by a user application as long as the operating system and user application are running.

When the application is rearming the application watchdog, an expiration can be caused by:

- Crash of the rearming application
- Hang or crash of the rearming thread in the application
- System hang

When the system watchdog is running, a system hang, or more specifically, the hang of the clock interrupt handler, causes an expiration.

The system watchdog mode is the default. If the application watchdog is not initialized, then the system watchdog mode is used.

The application mode enables you to:

- Configure the watchdog timer – Your applications running on the host can configure and use the watchdog timer, enabling you to detect fatal problems from applications and to recover automatically.
- Program Alarm3 – This capability enables you to generate this alarm in case of critical problems in your applications.

The `setupsc` command, an existing command for the ALOM CMT compatibility CLI (in ILOM), can be used to configure the recovery for the system watchdog *only*:

```
sc> setupsc
```

The recovery configuration for the application watchdog is set using input/output control codes (IOCTLs) that are issued to the `ntwdt` driver.

Watchdog Timer Limitations

The limitations of the watchdog timer mode include:

- In the case of the watchdog timer expiration detected by the system controller, the recovery is attempted only once. There are no further attempts of recovery if the first attempt fails to recover the domain.
- If the application watchdog is enabled and you break into the OpenBoot PROM by issuing the `break` command from the system controller's `sc>` prompt, the system controller automatically disables the watchdog timer.

Note – The system controller displays a console message as a reminder that the watchdog, from the system controller's perspective, is disabled.

However, when you re-enter the Solaris OS, the watchdog timer is still enabled from the Solaris OS perspective. To have both the system controller and the Solaris OS view the same watchdog state, you must use the watchdog application to either enable or disable the watchdog.

- If you perform a dynamic reconfiguration (DR) operation in which a system board containing kernel (permanent) memory is deleted, then you must disable the watchdog timer's application mode before the DR operation and enable it after the DR operation. This is required because Solaris software quiesces all

system IO and disables all interrupts during a memory-delete of permanent memory. As a result, system controller firmware and Solaris software can not communicate during the DR operation. Note that this limitation affects neither the dynamic addition of memory nor the deletion of a board not containing permanent memory. In those cases, the watchdog timer's application mode can run concurrently with the DR implementation.

You can execute the following command to locate the system boards that contain kernel (permanent) memory:

```
# cfgadm -lav | grep -i permanent
```

- If the Solaris Operating System hangs under the following conditions, the system controller firmware cannot detect the Solaris software hang:
 - Watchdog timer's application mode is set.
 - Watchdog timer is not enabled.
 - No rearming is done by the user.
- The watchdog timer provides partial boot monitoring. You can use the application watchdog to monitor a domain reboot.

However, domain booting is not monitored for:

- Bootup after a cold poweron.
- Recovery of a hung or failed domain.

In the case of a recovery of a hung or failed domain, a boot failure is not detected and no recovery attempts are made.

- The watchdog timer's application mode provides no monitoring for application startup. In application mode, if the application fails to start up, the failure is not detected and no recovery is provided.

Using the ntwdt Driver

To enable and control the watchdog's application mode, you must program the watchdog system using the LOMIOCDOGxxx IOCTLs, described in ["Understanding the User API" on page 78](#).

If the ntwdt driver, as opposed to the system controller, initiates a reset of the Solaris OS on application watchdog expiration, the value of the following property in the ntwdt driver's configuration file (ntwdt.conf) is used:

```
ntwdt-boottimeout="600";
```

In case of a panic, or an expiration of the application watchdog, the ntwdt driver reprograms the watchdog time-out to the value specified in the property.

Assign a value representing a duration that is longer than the time it takes to reboot and perform a crash dump. If the specified value is not large enough, the system controller resets the host if reset is enabled. Note that this reset by the system controller occurs only once.

Understanding the User API

The `ntwtdt` driver provides an application programming interface by using IOCTLs. You must open the `/dev/ntwtdt` device node before issuing the watchdog IOCTLs.

Note – Only a single instance of `open()` is allowed on `/dev/ntwtdt`. More than one instance of `open()` will generate the following error message: `EAGAIN` – The driver is busy, try again.

You can use the following IOCTLs with the watchdog timer:

- `LOMIOCDOGTIME`
- `LOMIOCDOGCTL`
- `LOMIOCDOGPAT`
- `LOMIOCDOGSTATE`
- `LOMIOCALSTATE`

Using the Watchdog Timer

Setting the Timeout Period

The `LOMIOCDOGTIME` IOCTL sets the timeout period of the watchdog. This IOCTL programs the watchdog hardware with the time specified in this IOCTL. You must set the timeout period (`LOMIOCDOGTIME`) before attempting to enable the watchdog timer (`LOMIOCDOGCTL`).

The argument is a pointer to an unsigned integer. This integer holds the new timeout period for the watchdog in multiples of 1 second. You can specify any timeout period in the range of 1 second to 180 minutes.

If the watchdog function is enabled, the time-out period is immediately reset so that the new value can take effect. An error (EINVAL) is displayed if the timeout period is less than 1 second or longer than 180 minutes.

Note – The LOMIOCDOGTIME is not intended for general-purpose use. Setting the watchdog time-out to too low a value might cause the system to receive a hardware reset if the watchdog and reset functions are enabled. If the timeout is set too low, the user application must be run with a higher priority (for example, as a real-time thread) and must be rearmed more often to avoid an unintentional expiration.

Enabling or Disabling the Watchdog

The LOMIOCDOGCTL IOCTL enables or disables the watchdog, and it enables or disables the reset capability. See [“Finding and Defining Data Structures” on page 80](#) for the correct values for the watchdog timer.

The argument is a pointer to the `lom_dogctl_t` structure. This structure is described in greater detail in [“Finding and Defining Data Structures” on page 80](#).

Use the `reset_enable` member to enable or disable the system reset function. Use the `dog_enable` member to enable or disable the watchdog function. An error (EINVAL) is displayed if the watchdog is disabled but reset is enabled.

Note – If LOMIOCDOGTIME has not been issued to set up the timeout period prior to this IOCTL, the watchdog is *not* enabled in the hardware.

Rearming the Watchdog

The LOMIOCDOGPAT IOCTL rearms, or pats, the watchdog so that the watchdog starts ticking from the beginning; that is, to the value specified by LOMIOCDOGTIME. This IOCTL requires no arguments. If the watchdog is enabled, this IOCTL must be used at regular intervals that are less than the watchdog timeout, or the watchdog expires.

Obtaining the State of the Watchdog Timer

The LOMIOCDOGSTATE IOCTL gets the state of the watchdog and reset functions, and retrieves the current time-out period for the watchdog. If LOMIOCDOGSTATE was never issued to set up the timeout period prior to this IOCTL, the watchdog is not enabled in the hardware.

The argument is a pointer to the `lom_dogstate_t` structure, which is described in greater detail in [“Finding and Defining Data Structures” on page 80](#). The structure members are used to hold the current states of the watchdog reset circuitry and current watchdog timeout period. This timeout period is not the time remaining before the watchdog is triggered.

The LOMIOCDOGSTATE IOCTL requires only that `open()` be successfully called. This IOCTL can be run any number of times after `open()` is called, and it does not require any other DOG IOCTLs to have been executed.

Finding and Defining Data Structures

All data structures and IOCTLs are defined in `lom_io.h`, which is available in the `SUNWl0mh` package.

The data structures for the watchdog timer are shown here:

- The watchdog and reset state data structure is as follows:

```
typedef struct {
    int reset_enable; /* reset enabled if non-zero */
    int dog_enable; /* watchdog enabled if non-zero */
    uint_t dog_timeout; /* Current watchdog timeout */
} lom_dogstate_t;
```

- The watchdog and reset control data structure is as follows:

```
typedef struct {
    int reset_enable; /* reset enabled if non-zero */
    int dog_enable; /* watchdog enabled if non-zero */
} lom_dogctl_t;
```

Example Watchdog Program

Following is a sample program for the watchdog timer.

EXAMPLE A-1 Example Watchdog Program (*Continued*)

```
#include <sys/types.h>
#include <fcntl.h>
#include <unistd.h>
#include <sys/stat.h>
#include <lom_io.h>

int main() {
    uint_t timeout = 30; /* 30 seconds */
    lom_dogctl_t dogctl;
    int fd;

    dogctl.reset_enable = 1;
    dogctl.dog_enable = 1;

    fd = open("/dev/ntwdt", O_EXCL);

    /* Set timeout */
    ioctl(fd, LOMIOCDOGTIME, (void *)&timeout);

    /* Enable watchdog */
    ioctl(fd, LOMIOCDOGCTL, (void *)&dogctl);

    /* Keep patting */
    while (1) {
        ioctl(fd, LOMIOCDOGPAT, NULL);
        sleep (5);
    }
    return (0);
}
```

Watchdog Timer Error Messages

TABLE A-1 describes watchdog timer error messages that might be displayed and

what they mean.

TABLE A-1 Watchdog Timer Error Messages

Error Message	Meaning
EAGAIN	An attempt was made to open more than one instance of <code>open()</code> on <code>/dev/ntwdt</code> .
EFAULT	A bad user-space address was specified.
EINVAL	A nonexistent control command was requested or invalid parameters were supplied.
EINTR	A thread awaiting a component state change was interrupted.
ENXIO	The driver is not installed in the system.

Alarm Library `libtsalarm`

The `libtsalarm` library program allows you to get or set the status of the alarms with the `tsalarm_get` and `tsalarm_set` functions. For more details on the alarm indicators, see [“Alarm Status Indicators” on page 35](#).

The following is an example of an application using the `libtsalarm` library.

EXAMPLE B-1 Application Using the libtsalarm Library (Continued)

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <sys/types.h>
#include <tsalarm.h>

void help(char *name) {
    printf("Syntax:  %s [get <type> | set <type> <state>]\n\n", name);
    printf("           type = { critical, major, minor, user }\n");
    printf("           state = { on, off }\n\n");

    exit(0);
}

int main(int argc, char **argv) {

    uint32_t alarm_type, alarm_state;

    if (argc < 3)
        help(argv[0]);

    if (strncmp(argv[2], "critical", 1) == 0)
        alarm_type = TSALARM_CRITICAL;
    else if (strncmp(argv[2], "major", 2) == 0)
        alarm_type = TSALARM_MAJOR;
    else if (strncmp(argv[2], "minor", 2) == 0)
        alarm_type = TSALARM_MINOR;
    else if (strncmp(argv[2], "user", 1) == 0)
        alarm_type = TSALARM_USER;
    else
        help(argv[0]);

    if (strncmp(argv[1], "get", 1) == 0) {
        tsalarm_get(alarm_type, &alarm_state);
        printf("alarm = %d\tstate = %d\n", alarm_type, alarm_state);
    }
    else if (strncmp(argv[1], "set", 1) == 0) {
        if (strncmp(argv[3], "on", 2) == 0)
            alarm_state = TSALARM_STATE_ON;
        else if (strncmp(argv[3], "off", 2) == 0)
            alarm_state = TSALARM_STATE_OFF;
        else
            help(argv[0]);

        tsalarm_set(alarm_type, alarm_state);
    }
    else {
```

```
        help(argv[0]);  
    }  
  
    return 0;  
}
```


OpenBoot Configuration Variables

TABLE C-1 describes the OpenBoot firmware configuration variables stored in non-volatile memory on the system. The OpenBoot configuration variables are printed here in the order in which they appear when you issue the following command:

```
-> show -o table -level all /SYS
```

TABLE C-1 OpenBoot Configuration Variables Stored on the System Configuration Card

Variable	Possible Values	Default Value	Description
ttya-rts-dtr-off	true, false	false	If true, operating system does not assert rts (request-to-send) and dtr (data-transfer-ready) on serial management port.
ttya-ignore-cd	true, false	true	If true, operating system ignores carrier-detect on serial management port.
keyboard-layout			
reboot-command			
security-mode	none, command, full	none	Firmware security level.
security-password	variable-name	none	Firmware security password if security-mode is not none (never displayed). Do not set this directly.
security-#badlogins	variable-name	none	Number of incorrect security password attempts.
verbosity	min, max	min	Sets verbosity mode.
pci-mem64?	true, false	false	

TABLE C-1 OpenBoot Configuration Variables Stored on the System Configuration Card *(Continued)*

Variable	Possible Values	Default Value	Description
diag-switch?	true, false	false	If true: <ul style="list-style-type: none"> • OpenBoot verbosity is set to maximum If false: <ul style="list-style-type: none"> • OpenBoot verbosity is set to minimum
local-mac-address?	true, false	true	If true, network drivers use their own MAC address, not the server MAC address.
fcode-debug?	true, false	false	If true, include name fields for plug-in device FCodes.
scsi-initiator-id	0-15	7	SCSI ID of the Serial Attached SCSI controller.
oem-logo		No default	
oem-logo?	true, false	false	If true, use custom OEM logo; otherwise, use the server manufacturer's logo.
oem-banner		No default	
oem-banner?	true, false	false	If true, use custom OEM banner.
ansi-terminal?	true, false	true	If true, enable ANSI terminal emulation.
screen-#columns	0-n	80	Sets number of columns on screen.
screen-#rows	0-n	34	Sets number of rows on screen.
ttya-mode	9600,8,n,1,-	9600,8,n,1,-	Serial management port (baud rate, bits, parity, stop, handshake). The serial management port only works at the default values.
output-device	virtual-console, screen	virtual-console	Power-on output device.
input-device	virtual-console, keyboard	virtual-console	Power-on input device.
auto-boot-on-error?	true, false	false	If true, boot automatically after system error.
load-base	0-n	16384	Address.
auto-boot?	true, false	true	If true, boot automatically after power on or reset.

TABLE C-1 OpenBoot Configuration Variables Stored on the System Configuration Card (Continued)

Variable	Possible Values	Default Value	Description
network-boot-arguments	[<i>protocol</i> ,] [<i>key</i> = <i>value</i> ,]	none	Arguments to be used by the PROM for network booting. Defaults to an empty string. <code>network-boot-arguments</code> can be used to specify the boot protocol (RARP/DHCP) to be used and a range of system knowledge to be used in the process. For further information, see the <code>eeprom (1M)</code> man page or your Solaris reference manual.
boot-command	<i>variable-name</i>	boot	Action following a boot command.
boot-file			
boot-device	/pci@0/pci@0/pci@2/scsi@...	disk net	
multipath-boot?	true, false	false	
boot-device-index	0-n	0	
use-nvramrc?	true, false	false	If true, execute commands in NVRAMRC during server startup.
nvramrc	<i>variable-name</i>	none	Command script to execute if <code>use-nvramrc?</code> is true.
error-reset-recovery	boot, sync, none	boot	Command to execute following a system reset generated by an error.

Index

Symbols

/etc/remote file, 14

modifying, 14

-> commands

reset /SYS, 39

set /HOST/bootmode, 39

set /SYS/LOCATE, 37, 38

show /SYS/LOCATE, 38

-> prompt

about, 20, 30

accessing from network management port, 21

accessing from serial management port, 21

multiple sessions, 21

system console escape sequence (#.), 21

system console, switching between, 18

ways to access, 21

A

Activity (disk drive LED), 69

activity indicator, 35

alarm

status indicators, 36

alphanumeric terminal

accessing system console from, 15

setting baud rate, 16

auto-boot (OpenBoot configuration variable), 22, 41

automatic system recovery (ASR)

about, 40

commands, 43

disabling, 44

enabling, 43

obtaining recovery information, 45

B

Break key (alphanumeric terminal), 25

C

cables, keyboard and mouse, 17

cfgadm (Solaris command), 67

cfgadm install_device (Solaris command),
cautions against using, 68

cfgadm remove_device (Solaris command),
cautions against using, 68

Cisco AS2511-RJ Terminal Server, connecting, 10

command prompts, explained, 19

communicating with the system

about, 1

options, table, 2

console configuration, connection alternatives
explained, 6

critical alarm, 36

D

default system console configuration, 4

device reconfiguration, manual, 47

device unconfiguration, manual, 45

device, disabling, 46

device, enabling, 47

disabling a device, 46

disk configuration

RAID 0, 52

RAID 1, 53

- disk drives
 - LEDs
 - Activity, 69
 - OK-to-Remove, 68, 69
 - logical device names, table, 54
- disk hot-plug
 - mirrored disk, 65
 - nonmirrored disk, 66
- disk slot number, reference, 55
- disk volumes
 - about, 51
 - deleting, 64
- dtterm (Solaris utility), 14
- Dynamic Host Configuration Protocol (DHCP)
 - client on network management port, 10

E

- enabling a device, 47
- environmental information, viewing, 31
- error handling, summary, 41
- escape sequence (#.), service processor, 21

G

- graceful system halt, 23, 25
- graphics monitor
 - accessing system console from, 16
 - connecting to PCI graphics card, 17
 - restrictions against using for initial setup, 16
 - restrictions against using to view POST output, 16

H

- halt, gracefully, advantages of, 23, 25
- hardware disk mirror
 - about, 54
 - hot-plug operation, 65
- hardware disk mirrored volume
 - checking the status of, 58
- hardware disk stripe
 - about, 52
- hardware disk striped volume
 - checking the status of, 61
- hardware disk striping, about, 52
- hot-plug operation
 - nonmirrored disk drive, 66
 - on hardware disk mirror, 65

I

- ILOM commands
 - reset /SYS, 25
 - show, 10
- ILOM, *See* Integrated Lights Out Manager (ILOM)
- init (Solaris command), 23, 25
- input-device (OpenBoot configuration variable), 17, 27
- Integrated Lights Out Manager (ILOM)
 - > prompt, *See* -> prompt
 - commands, *See* -> prompt
 - escape sequence (#.), 21
 - logging in, 30
 - multiple connections to, 21

K

- keyboard, attaching, 17

L

- LEDs, 31
 - Activity (disk drive LED), 69
 - alarm status, 33
 - critical, 36
 - major, 36
 - minor, 37
 - user, 37
 - OK-to-Remove (disk drive LED), 68, 69
 - rear panel, 5
 - server status, 33
- LEDs, Locator (system status LED), 37
- Locator (system status LED)
 - controlling from -> prompt, 37, 38
- Locator (system status LED), controlling, 37
- locator indicator, 35
- logging in to the Integrated Lights Out Manager (ILOM), 30
- logical device name (disk drive), reference, 54

M

- major alarm, 36
- manual device reconfiguration, 47
- manual device unconfiguration, 45
- manual system reset, 25, 26
- minor alarm, 37
- monitor, attaching, 16
- multiple ILOM sessions, 21

N

- network management port (NET MGT)
 - activating, 8
- nonmirrored disk hot-plug operation, 66
- normally
 - closed (NC) relay state, 37
 - open (NO) relay state, 37
- ntwdt driver, 77

O

- ok prompt
 - about, 22
 - accessing with Break key, 24
 - accessing with graceful system shutdown, 23
 - accessing with ILOM set /HOST
 - break_action=break command, 24
 - accessing with ILOM set /HOST
 - send_break_action=break command, 23
 - accessing with manual system reset, 23, 25
 - suspension of Solaris operating system, 25
 - ways to access, 23
- OK-to-Remove (disk drive LED), 68, 69
- OpenBoot commands
 - probe-ide, 24
 - probe-scsi-all, 24
 - reset-all, 17
 - setenv, 17
- OpenBoot configuration variables
 - auto-boot, 22, 41
 - described, table, 87
 - input-device, 17, 27
 - output-device, 17, 27
 - system console settings, 26
- OpenBoot emergency procedures
 - performing, 38
 - USB keyboard commands, 38
- OpenBoot firmware
 - scenarios for control, 22
- output-device (OpenBoot configuration variable), 17, 27

P

- parity, 16
- patch panel, terminal server connection, 11
- PCI graphics card
 - configuring to access system console, 16
 - connecting graphics monitor to, 17

- frame buffers, 16

- physical device name (disk drive), 54
- probe-ide (OpenBoot command), 24
- probe-scsi-all (OpenBoot command), 24

R

- RAID (redundant array of independent disks), ix, 51
- RAID 0 (striping), 52
- RAID 1 (mirroring), 53
- raidctl (Solaris command), 55 to 66
- rear panel
 - LEDs, 5
- relay state
 - normally closed (NC), 37
 - normally open (NO), 37
- reset
 - manual system, 25, 26
 - scenarios, 42
- reset /SYS (ILOM command), 25
- reset-all (OpenBoot command), 17
- run levels
 - explained, 22
 - ok prompt and, 22

S

- SER MGT, *See* serial management port
- serial management port (SER MGT)
 - acceptable console device connections, 4
 - as default communication port on initial startup, 2
 - configuration parameters, 7
 - default system console configuration, 4
 - using, 7
- service required indicator, 35
- set /HOST/bootmode (-> command), 39
- set /SYS/LOCATE (-> command), 38
- setenv (OpenBoot command), 17
- show (ILOM CMT command), 10
- shutdown (Solaris command), 23, 25
- Solaris commands
 - cfgadm, 67
 - cfgadm install_device, cautions against using, 68
 - cfgadm remove_device, cautions against using, 68

- init, 23, 25
- raidctl, 55 to 66
- shutdown, 23, 25
- tip, 12, 14
- uadmin, 23
- uname, 15
- uname -r, 14
- status indicators, 31
 - alarm, 33, 36
 - critical, 36
 - major, 36
 - minor, 37
 - user, 37
 - server, 33
- Stop-D (USB keyboard functionality), 40
- Stop-F (USB keyboard functionality), 40
- Stop-N (USB keyboard functionality), 38
- system console
 - > prompt, switching between, 18
 - accessing with alphanumeric terminal, 15
 - accessing with graphics monitor, 16
 - accessing with terminal server, 2, 10
 - accessing with tip connection, 12
 - alphanumeric terminal connection, 2, 15
 - alternate configurations, 6
 - configuring local graphics monitor to access, 16
 - connection using graphics monitor, 7
 - default configuration explained, 2, 4
 - default connections, 4
 - defined, 1
 - Ethernet attachment through network
 - management port, 2
 - graphics monitor connection, 2, 7
 - multiple view sessions, 21
 - setting OpenBoot configuration variables for, 26
- system reset scenarios, 42
- system status LEDs
 - Locator, 37
- system status LEDs, Locator, 37

T

- terminal server
 - accessing system console from, 4, 10
 - connection through patch panel, 11
 - pinouts for crossover cable, 12
- tip (Solaris command), 14
- tip connection

- accessing system console, 12
- accessing terminal server, 12
- tip connection
 - accessing system console, 12

U

- uadmin (Solaris command), 23
- uname (Solaris command), 15
- uname -r (Solaris command), 14
- user alarm, 37

W

- watchdog timer
 - APIs, 78
 - application mode, 75
 - data structures, 80
 - disabling, 79
 - enabling, 79
 - error messages, 81
 - example program, 80
 - getting state of, 80
 - IOCTLs, 78
 - limitations, 76
 - rearming, 79
 - setting timeout period, 78